

PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION  
International Bureau




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<b>(51) International Patent Classification <sup>6</sup> :</b> <b>H04N 7/58</b>		<b>A1</b>	<b>(11) International Publication Number:</b> <b>WO 97/15149</b>
			<b>(43) International Publication Date:</b> 24 April 1997 (24.04.97)
<b>(21) International Application Number:</b> PCT/IB96/01074		<b>(81) Designated States:</b> CN, JP, KR, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).	
<b>(22) International Filing Date:</b> 11 October 1996 (11.10.96)			
<b>(30) Priority Data:</b> 95202811.6 18 October 1995 (18.10.95) EP <b>(34) Countries for which the regional or international application was filed:</b> AT et al.		<b>Published</b> <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>	
<b>(71) Applicant:</b> PHILIPS ELECTRONICS N.V. [NL/NL]; Groenewoudseweg 1, NL-5621 BA Eindhoven (NL).			
<b>(71) Applicant (for SE only):</b> PHILIPS NORDEN AB [SE/SE]; Kottbygatan 7, Kista, S-164 85 Stockholm (SE).			
<b>(72) Inventors:</b> ELBERS, Gerard, J.; Donk 154, NL-4824 CZ Breda (NL). BOUFFIOUX, Alain, R., P.; Froidbermont 8B, B-4877 Olne (BE).			
<b>(74) Agent:</b> STRIJLAND, Wilfred; Internationaal Octrooibureau B.V., P.O. Box 220, NL-5600 AE Eindhoven (NL).			

**(54) Title:** METHOD FOR MAKING A MULTIMEDIA APPLICATION EXECUTABLE ON HARDWARE PLATFORMS WITH VARIOUS DIFFERENT RESOURCE LEVELS, A PHYSICAL RECORD CONTAINING SUCH APPLICATION, AND AN APPARATUS FOR EXECUTING SUCH APPLICATION

**(57) Abstract**

Multimedia application information is stored for later executability on different platforms, such as Interactive Multimedia Distribution Format (IMDF) complying platforms. The application comprises a plurality of parallel information streams. Execution scalability is enabled by assigning to the streams a single string of quality level items that along the string indicate stepwise increasing quality levels. Each item specifies an associated decoding quality level for a particular stream.



**FOR THE PURPOSES OF INFORMATION ONLY**

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AM	Armenia	GB	United Kingdom	MW	Malawi
AT	Austria	GE	Georgia	MX	Mexico
AU	Australia	GN	Guinea	NE	Niger
BB	Barbados	GR	Greece	NL	Netherlands
BE	Belgium	HU	Hungary	NO	Norway
BF	Burkina Faso	IE	Ireland	NZ	New Zealand
BG	Bulgaria	IT	Italy	PL	Poland
BJ	Benin	JP	Japan	PT	Portugal
BR	Brazil	KE	Kenya	RO	Romania
BY	Belarus	KG	Kyrgyzstan	RU	Russian Federation
CA	Canada	KP	Democratic People's Republic of Korea	SD	Sudan
CF	Central African Republic	KR	Republic of Korea	SE	Sweden
CG	Congo	KZ	Kazakhstan	SG	Singapore
CH	Switzerland	LI	Liechtenstein	SI	Slovenia
CI	Côte d'Ivoire	LK	Sri Lanka	SK	Slovakia
CM	Cameroon	LR	Liberia	SN	Senegal
CN	China	LT	Lithuania	SZ	Swaziland
CS	Czechoslovakia	LU	Luxembourg	TD	Chad
CZ	Czech Republic	LV	Latvia	TG	Togo
DE	Germany	MC	Monaco	TJ	Tajikistan
DK	Denmark	MD	Republic of Moldova	TT	Trinidad and Tobago
EE	Estonia	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	UG	Uganda
FI	Finland	MN	Mongolia	US	United States of America
FR	France	MR	Mauritania	UZ	Uzbekistan
GA	Gabon			VN	Viet Nam

Method for making a multimedia application executable on hardware platforms with various different resource levels, a physical record containing such application, and an apparatus for executing such application.

## DEFINITIONS

The following definitions are given for better understanding:

Resource	quantitative and/or qualitative properties of the platform such as MIPS of the CPU, memory size in Mbytes, bandwidth of bus in Mbits/sec.
Resource requirement	necessary amount of such resources
Performance	speed of a computer or CPU
Quality level	level of functioning of a decoder, which has been defined independently of the platform

## BACKGROUND TO THE INVENTION

The invention relates to a method for storing multimedia application information for execution on different platforms, said information comprising a plurality of information streams for presentation in parallel to a user. Multimedia application programs or applications should be executable on different platforms. Such platform must satisfy minimum requirements in terms of resources, for being able to guarantee that the application in question runs thereon correctly. In various cases the platform has more resources than the minimum. Advantageously, the application should be made scalable, to attain optimum quality level, given the resources of the actually available platform.

## SUMMARY TO THE INVENTION

In consequence, amongst other things it is an object of the present invention to provide a method for storing multimedia application information, wherein such scalability is realized in a straightforward manner. Now, according to one of its aspects, the invention is characterized in that for enabling execution scalability of said application, said method collectively assigns to said streams a single string of quality level items that along said string indicate stepwise increasing quality levels, in that each item specifies an associated decoding quality level for a particular one of said streams. Their relative positions in the string give the priorities among the various streams, and also between the various quality

levels for each single stream. Now first, for all streams together, specified minimum platform resources are necessary for the platform to executing the application or not.

Moreover, if additional resources are present, assigning thereof to the various streams on the basis of relative priorities is straightforward. This organization allows the platform to easily

5 choose optimum quality level with respect to available resources.

By itself, multimedia executability on different platforms has been considered in EP A 731 971, corresponding US Patent Application 08/533,030 (PHN 15058) to the present assignee, and herein incorporated by reference. In the reference, the multimedia application itself has a multiplatform interpreter, and each actual platform is an instance of an abstract  
10 machine. The present invention may be used in such more complex environment, but it is already useful if platform resources differ only in quantitative performance.

In the organization according to the invention, the platform or player knows the resource requirements for each quality level of each decoder. These have been stored in the resource requirement table of the platform (cf. Figure 5). Neither the data streams, nor the  
15 application itself need to know these, because generally, at configuring the data streams and writing the application, the target platform is not known yet. The data streams comprise the input data of all associated decoders. In general, no further information is required for realizing the organization of the present invention. The application comprises a model, realized as a real or notional string of items, which describes the incremental model or  
20 preference path of the application in view of respective decoder levels.

The resource requirements of a particular platform are fixed for each decoder and each associated quality level when configuring the platform. These are stored in the platform. The producing of the application implies the joining of the data streams in a real-time file, using the so-called discbuilder. Furthermore, an incremental model of the  
25 application is built in the application, with one or more calls thereto. Generally, this model is separate from the data streams.

Advantageously, at least one such stream is associated to at least two non-contiguous items of said string. This means that when stream A gets a particular quality level, first the quality of stream B will be improved, before further improvement of stream A  
30 can be allowed.

Advantageously, the quality level items are assigned to said streams in a dynamic manner along said application when stored. The writer of an application can thereby decide that the preference amongst the various decoders may vary during the executing of the application. During the execution proper on a particular platform, this may influence the

quality levels actually used, or not.

The invention also relates to a unitary physical record carrier containing such application and to a platform system for playing such multimedia application information, wherein the platform will automatically select appropriate quality levels for the various streams. Further advantageous aspects of the invention are recited in dependent Claims.

## BRIEF DESCRIPTION OF THE DRAWING

These and further aspects and advantages of the invention will be described more in detail with reference to preferred embodiments disclosed hereinafter, and in particular with reference to the appended Figures that show:

Figure 1 a flow chart for storing an application according to the invention;

Figure 2 a unitary record carrier so formed;

Figure 3 a platform system for playing the record;

Figure 4 a quality of service management diagram;

Figure 5 a platform resource requirement table;

Figure 6 a preference path diagram.

## GENERAL CONSIDERATIONS

In its most complete form, a multimedia and in particular, an IMDF application may contain two types of processing:

- Stream Processing (SP) of audio-visual information streams, such as Video (MPEG), 2D- or 3D-graphics, audio (MPEG, MIDI, PCM, ADPCM), which streams must be processed in a time-conforming manner at the normal rendering speed a human user would expect.
- Non-Stream processing (NSP) that is not directly related to such time-conforming audio-visual streams. For simplicity of wording, the informations relating thereto on the physical record carrier are nevertheless also called streams. This processing type can be subdivided into:
  - (a) Real-time constrained NSP, such as simulation calculations, e.g. a flight simulator
  - (b) Non-real-time constrained NSP, such as a search in a data base or encyclopedia, for presenting stationary information.

At its low end, each application requires a platform with minimum resources; these may be defined in an applicable standard. On such minimum platform only a limited number of information streams and computations can be processed in parallel. A more powerful platform will in principle allow for a higher quality level compared to the mini-

5 mum. Examples of such resources are the recurrence frequency of CPU processing cycles, memory size, bus transfer capacity, and I/O-facilities. In certain circumstances the MIPS figure indicating the raw instruction processing speed may need recalculation to MOPS, that is the operation speed as pertaining to the stream in question. Also operational overhead, such as for loading of the whole or part of the application, may need to be taken into account. In case of a server or other multi-process environment, available resources may have to be shared among various applications that run concurrently.

For stream processing (SP), the resource management is done in three steps:

- defining objective quality levels for all feasible decoders;
- 10 - selecting an actual quality level for a particular decoder;
- controlling decoder resources at a selected quality level. Herein, decoders in fact are decoder programs that run in parallel on hardware that may be shared or not among various decoders.

15 In a first approach, non-stream processing (NSP) is considered to be fully determined by the application and non-scalable, as the platform cannot foresee what the NSP will be doing. This means that for NSP resource requirements must be known beforehand in the platform. The amount of allocated resources must be sufficient so that the application will always meet its deadlines if real-time constrained. If not real-time constrained, the allocated resources should be sufficient to avoid any long-term build-up of unfinished business. This  
20 generally means that the resources for NSP must be assigned as a whole. Subsequently, any remaining resources can be assigned to SP.

## DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

Figure 1 shows a flow chart for storing an application according to the  
25 invention. The operations according to this flow chart can be executed directly after the general writing of the application. In block 20 the unprocessed information of the application is received, as pertaining to a particular time interval or scene. Generally, the scene size has been determined by the developer of the application. In so doing, the scene has been usually built up from things like music or movie intervals, stills, graphic elements and others. In the  
30 information, the data streams have usually been joined to an interleaved and multiplexed data file. Sufficient processing power is available for executing the application at its highest quality level. During execution of the application, often the processing is linear in-time, in that the various streams are presented to a user in mutual synchronism, such as in the way of a movie. However, this behaviour need not be uniform in time, where at one time an image

may freeze, while accompanying music goes on. In an interactive situation, the relative behaviours may be still more non-uniform. In block 22 the first unprocessed item of the application is accessed, in view of therefor determining its string of quality level items.

In block 28, the minimum quality level for executing the (part of the) applica-  
5 tion is assessed. This decision may be taken by the writer of the application. Next, the scene heading is now provided the minimally required quality levels of all decoders as a first item of the string of quality level items. Inasmuch as this first item may pertain to more than one decoder, it may be a composite item. For being able to attain automatic scalability, subse-  
10 quent items of the string must now be arranged so that a stepwise increasing quality can be indicated by going along the items of the string. The operation in block 28 can be done either with the isolated scene, or in the context of preceding and/or succeeding scenes. Next, in block 30, exactly one of the streams which allows a plurality of different quality levels is taken at the next higher quality level, and the scene is executed again. This is repeated for each the streams that have incrementable quality levels, which is symbolized by the looping  
15 line 31. When the scene has been executed with all possibilities for single-quality increment, in block 32 the best performance thereamongst is chosen. This is generally a subjective action that is done by a designer person. This determines the next item of the string. As long as further quality improvement is possible (34), the system next goes back to block 30. If all streams have been executed at their respective top quality levels, the application scene is  
20 loaded (36). The ultimate result of the invention is a unitary carrier that is manufactured in standard manner for media of this kind. In block 38 the processing system is relinquished. For simplicity, only the loading of a single scene has been considered. However, one may decide to execute the flow chart of Figure 1 again for the next scene.

Figure 2 symbolizes a unitary record carrier so formed. The carrier contains  
25 eight streams shown as separate columns. For physical storage, generally a serializing operation is effected, such as for Compact-Disc Interactive; this can produce a real-time multiplexed file. Each streams accommodates a block 32 that contains an identifier, length, and various other references. The parallel character of the streams resides in the fact that they can be accessed for replay in an interleaved manner. For brevity, actual organization of  
30 physical storage has not been addressed. An exemplary solution has been disclosed in copending EP Patent Application 95202623.5, corresponding United States Patent Application 08/..... (PHN 15480) to the present assignee and herein incorporated by reference. In such interleaving, the player keeps track of the stream that is actually being accessed. As recited supra, the string of quality level items which indicates the stepwise increasing quality

levels is stored in a scene header indicated by X in the Figure.

Figure 3 shows a platform system for playing the record carrier. The disc-shaped record 40 is accessed by read arrangement 42 under control of positioning signals on line 44. The player proper 60 has user interface 50 for inputting user commands and detector block 46 for detecting the string of quality level items recited supra. This block also knows the required resources for the various quality levels of the various decoders, as pertaining to the platform in question. Therefore it may select the actually attainable level of quality, in that certain requirements can be met, but others would overtax the platform. This will hereinafter be discussed further. Central processing device 48 processes the streams in question in accordance with the above selection. Such processing may further be under control by user commands on interface 50. In principle, the user may overrule the quality levels selected by the platform. For example, a user may prefer audio representation above video, and therefor by preference allocate more resources to audio processing than according to the relative priorities assigned thereto by the designer of the application. Device 48 outputs various user information signals to various user output facilities 52, 54, 56, 58, that may contain audio channels, video channels, and data channels. Various ones of these may be collocated, such as the displaying of video and animation in parallel on the same screen. The processing also causes positioning on line 44 for read arrangement 42.

Hereinafter, resource management will be discussed more in detail. The processing device 48 includes various facilities, such as memory, bus, and I/O features. The information thereof is known to detector block 46, which also knows the amounts of resources necessary for the various quality levels. This allows device 46 to choose an optimum quality level that can be realized with the resources actually present, as pertaining to the actual application. If an appropriate level of performance has been chosen, the necessary resources are allocated, and the application may be executed. If the resources are however insufficient such as defined in terms of bandwidth, delay, error rate, or other, the player will inform the user and the execution of the application will remain blocked. Upon meeting a subsequent scene, the resource allocation may have to recommence. If the resource requirements are dynamically varying along the application, the application may signal minimum overall quality level required, so that the overall executability of the application may be decided by the platform. Subsequently, instantaneous quality of service may be non-uniform in time.

In this context, Figure 4 shows a quality of service management diagram. Application 70 requests a particular quality of service to manager 72, which returns a



contract proposal. If accepted, the manager sends the parameters to the processing. The latter does runtime resource allocation (74), accounting and policing (76), and controls hardware elements 78-84. Resources so allocated render their services to application 70.

Figure 5 shows an exemplary resource requirement table that is present in the platform, inasmuch as the various streams, and consequently, the associated resource requirements have largely been standardized. Moreover, the application itself would not know the various resources present in the platform.

For MPEG video, the following objective quality levels for decoders may apply:

- 1. Black and white/half frame rate
- 2. Black and white/full frame rate
- 3. Full colour/half frame rate
- 4. Full colour/full frame rate

For MPEG audio, the following objective quality levels for decoders may apply:

- 1. Mono
- 2. Stereo
- 3. Dolby surround
- 4. 5+1 surround (5 standard channels plus one for special effects)

For 3-D graphics, the following objective quality levels for decoders may apply:

- 1. 5000 triangles or polygons per frame
- 2-4. 10k, 15k, and 20k triangles per frame, respectively.

In Figure 5, the columns list the resources required by these various quality levels from known facilities in the platform. The requirements as listed may pertain to averages, average plus a safety margin such as 20%, or to worst case. Herein, resources 1-4 may be first processor, second processor, Memory, and I/O facilities, respectively. Each resource is an abstract entity characterized by a maximum performance value.

In case the available resources extend beyond the bare minimum necessary for executing a particular application, the platform must decide as to which decoder should perform at what quality level. For this deciding, it is necessary to know the relative importance of a particular stream to the application, and also the necessary resources for its execution; the first knowledge is brought about by the application, inasmuch as it could widely vary between applications; the second is present in the platform itself. A straightfor-

ward way is to describe the executability levels in an incrementing manner. The term **decoder** is used for the actual decoding software package. For example:

Application: create MPEG2 video decoder 1 at quality level 1

create MPEG2 audio decoder 1 at quality level 1

5 raise MPEG2 audio decoder 1 to quality level 2

raise MPEG2 video decoder 1 to quality level 2

raise MPEG2 audio decoder 1 to quality level 3

player 1 \_\_\_\_\_

raise MPEG2 video decoder 1 to quality level 3

10 create MIDI-16 decoder 1 at quality level 1

player 2 \_\_\_\_\_

raise MIDI-16 decoder 1 to quality level 2

raise MIDI-16 decoder 1 to quality level 3

15 As shown in this table, for full functionality of the application, all three decoders should operate at quality level 3. For each of two exemplary available platforms, the actual performance is lower, and is given by the maximum level attained above the interrupted line for each of the decoders in question. For example, a relatively simple player has no MIDI, video at level 2, and audio at level 3. A relatively higher performing player has MIDI at level 1, and both audio and video at level 3. Obviously, the contribution by  
20 MIDI to the application is judged to be of relatively less importance. In another application, the relative importances of the streams could be different. For example, MIDI could be rated at approximately the same level of relevance for quality as audio. A different player could have a so-called 'sound blasters' card on board, which causes MIDI decoding taking place  
25 outside the central processor facilities.

In this context, Figure 6 shows a two-dimensional preference path diagram pertaining to a particular application, that has one dimension for each separate information stream of the application in question. For simplicity, a two-stream application has been taken in this example. The preference path is indicated as a solid line, and each cross therealong  
30 represents a sensible set of resources. Successive crosses indicate successive quality levels of playing the application. The preference path starts at the top left hand corner indicating zero resources. Subsequent crosses along the solid line correspond to successive rows in the same manner as given for the three-stream application considered earlier. Now, a particular cross (such as the one encircled) indicates minimum resources requirements for executing the

application on the platform in question. The platform tries to attain the best quality level, as represented by the farthest cross along the line of crosses.

Moreover, the dotted line indicates the full resources of the platform. The last cross not beyond this line indicates the maximum quality available by this particular platform. If the application has many streams in parallel, the matrix has as many dimensions, and moreover need not be static. In certain situations, the platform need not to keep strictly along the line of the diagram. For example, if MPEG-3 audio would just overtax the platform, the latter could still try whether MPEG-3 video (later in the sequence of steadily improving quality levels) could be feasibly executed by the platform. If yes, the latter may be included in the execution.

For the application, the specifying of the relative importance sequence may be a laborious task, especially if the number of streams is large and the assignment of relative importances is dynamic. Note that the number of streams may indeed be large, inasmuch as the number of languages may increase, the audio may be standard or karaoke, and also video may presented in more versions in parallel. To simplify this, a relative importance is assigned to each decoder, and these are mapped on an elementary model. Such is done as follows:

- The importance of a decoder is its minimally allowed quality level; this will corresponds to its first appearance in the incremental description.
- The quality of each decoder is incremented in the order in which they are opened.
- The quality of each decoder is incremented till it reaches maximum quality.

For example, an importance scheme with highest quality for each decoder at 3, and given as:

```

Decoder1 = OpenDecoder (MPEG2, Video,2)
Decoder2 = OpenDecoder (MPEG2, Audio,2)
Decoder3 = OpenDecoder (MIDI-16,1)
will give the following incremental description:
create MPEG2 video decoder 1 at quality level 2;
create MPEG2 audio decoder 1 at quality level 2;
create MIDI-16 decoder 1 at quality level 1;
raise MPEG2 video decoder 1 to quality level 3;
raise MPEG2 audio decoder 1 to quality level 3;
raise MIDI-16 decoder 1 to quality level 2;
raise MIDI-16 decoder 1 at quality level 3.
```

To select the quality level of each decoder, we start at the beginning of the incremental model. The decoders are created at a certain quality level and each decoder quality level is incremented following the preference path described in the incremental model. For each decoder that is created and for each quality level that is incremented, the  
5 resource requirements are looked up in the resource requirement table and added up to get the total resource requirements of the application. If the requirements for one of the resources exceeds its maximum value (i.e. 100%), no more decoder quality is implemented. This is the end of the preference path of this platform for the actual application or scene.

CLAIMS:

1. A method for storing multimedia application information for execution on different platforms, said information comprising a plurality of information streams for presentation in parallel to a user,  
characterized in that for enabling execution scalability of said application, said  
5 method collectively assigns to said streams a single string of quality level items that along said string indicate stepwise increasing quality levels, in that each item specifies an associated decoding quality level for a particular one of said streams.
2. A method as claimed in Claim 1, wherein at least one such stream is associated to at least two non-contiguous items of said string, that are separated by at least one item  
10 pertaining to another stream.
3. A method as claimed in Claims 1 or 2, wherein the quality level items are assigned to said streams in a dynamic manner along said application when stored.
4. A method as claimed in Claims 1, 2 or 3, for a number of streams that is at least equal to three and to each of which are assigned at least two such quality levels.
- 15 5. A unitary physical record carrier containing a comprehensive multimedia application stored according to the method as claimed in any of Claims 1 to 4,  
characterized by for enabling execution scalability of said application, said application has collectively assigned to said streams a single string of quality level items that along said string indicate a stepwise increasing quality levels, in that each item specifies an  
20 associated amount of associated decoding quality level for a particular one of said streams.
6. A record as claimed in Claim 5, wherein at least one such stream is associated to at least two non-contiguous items of said string, that are separated by at least one item pertaining to another stream.
7. A record as claimed in Claims 5 or 6, and comprising a plurality of said  
25 strings, each pertaining to a particular part of said application, and wherein the quality level items are assigned to said streams in a dynamic manner in said respective strings.
8. A record as claimed in Claims 5, 6 or 7, for a number of streams that is at least equal to three and to each of which are assigned at least two such quality levels.
9. A platform system for playing multimedia application information from a

storage medium, which application is executable on a range of different platforms and comprises a plurality of parallel information streams,

characterized in that for enabling execution scalability along said range, wherein said application has collectively assigned to said streams a single string of quality level items  
5 that along said string indicate stepwise increasing quality levels, in that each item specifies an associated quality level for a particular one of said streams,

said platform has interrogating means for interrogating said items along said string and accumulating all platform resources required for a particular quality level item, and comparing means fed by said interrogating means for comparing required platform  
10 resources to actual platform resources, and selecting means fed by said comparing means for selecting an optimum quality level as providable by the platform for the application in question.

10. A platform system as claimed in Claim 9, and having search means for searching additional string items beyond an item representing matching required platform  
15 resources to maximum available platform resources, whilst passing beyond an item overtaxing platform resources and pertaining to a first information stream, in favour of an item pertaining to a second information stream and commensurate with available platform resources, therewith improving said performance quality.

1/4

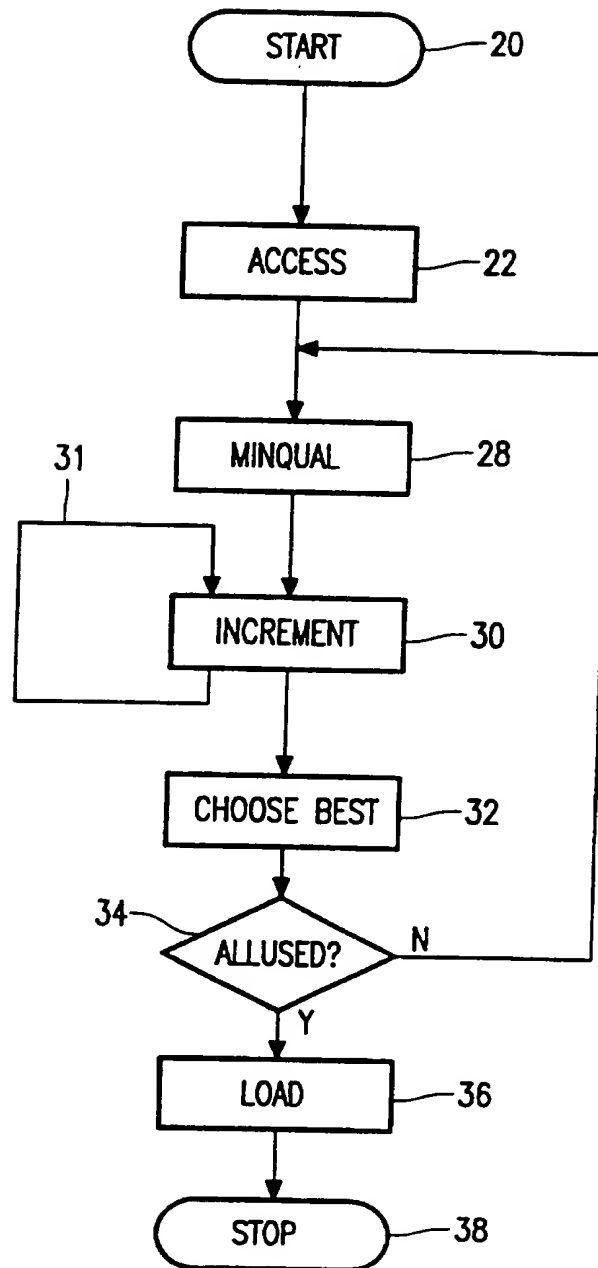


FIG. 1

2/4




FIG. 2




FIG. 3



3/4




FIG. 4




FIG. 6

4/4

An IMDF platform will have the following resource requirement table in memory

Decoder type	Resource 1	Resource 2	Resource 3	Resource 4
MPEG-video level 1	5%	20%	2 MB	0%
MPEG-video level 2	10%	20%	2 MB	0%
MPEG-video level 3	10%	20%	2 MB	0%
MPEG-video level 4	20%	20%	4 MB	0%
MPEG-audio level 1	...			
MPEG-audio level 2	...			
MPEG-audio level 3				
MPEG-audio level 4				
3D-Graphics level 1				
3D-Graphics level 1				
3D-Graphics level 1				
3D-Graphics level 1	25%	0%	4 MB	30%

FIG. 5

1  
INTERNATIONAL SEARCH REPORT

International application No.  
PCT/IB 96/01074

**A. CLASSIFICATION OF SUBJECT MATTER**

**IPC6: H04N 7/58**

According to International Patent Classification (IPC) or to both national classification and IPC

**B. FIELDS SEARCHED**

Minimum documentation searched (classification system followed by classification symbols)

**IPC6: H04N**

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

**SE,DK,FI,NO classes as above**

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5196933 A (JEAN-PIERRE HENOT), 23 March 1993 (23.03.93), column 2, line 15 - column 4, line 25 --	1-10
A	US 5216503 A (WOO H. PAIK), 1 June 1993 (01.06.93), column 1, line 62 - column 3, line 24 --	1-10
A	US 5363138 A (HAYASHI ET AL), 8 November 1994 (08.11.94), column 1, line 49 - column 3, line 41 --	1-10
A	EP 0643537 A2 (PHILIPS ELECTRONICS N.V.), 6 Sept 1994 (06.09.94), page 2, line 1 - page 3, line 2 -- -----	1-10

☐ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

\* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance: the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance: the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

Date of mailing of the international search report

**4 March 1997**

**06.03.1997**

Name and mailing address of the ISA/

Swedish Patent Office  
Box 5055, S-102 42 STOCKHOLM

Facsimile No. +46 8 666 02 86

Authorized officer

**KENNETH AHRENGART**

Telephone No. +46 8 782 25 00

**INTERNATIONAL SEARCH REPORT**  
Information on patent family members

03/02/97

International application No.  
PCT/IB 96/01074

Patent document cited in search report		Publication date	Patent family member(s)		Publication date
US-A-	5196933	23/03/93	DE-D,T-	69119634	23/01/97
			EP-A,B-	0448491	25/09/91
			FR-A,B-	2660139	27/09/91
US-A-	5216503	01/06/93	AU-B-	650219	09/06/94
			AU-A-	2995592	01/07/93
			CA-A-	2084178	25/06/93
			EP-A-	0550843	14/07/93
			JP-A-	6086261	25/03/94
US-A-	5363138	08/11/94	JP-A-	5122538	18/05/93
EP-A2-	0643537	06/09/94	AU-A-	7290294	23/03/95
			BE-A-	1007490	11/07/95
			CA-A-	2131573	11/03/95
			CN-A-	1111874	15/11/95
			JP-A-	7111649	25/04/95

(D3)

Internet Engineering Task Force  
INTERNET-DRAFT  
draft-christ-rtsp-mpeg4-00.txt

P. Christ, Ch. Guillemot, S. Wesner  
Univ. Stuttgart - RUS/ INRIA  
November 16, 1998  
Expires: May 15, 1999

RTSP-based Stream Control in MPEG-4

Status of this Memo

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months. Internet-Drafts may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet-Drafts as reference material or to cite them other than as a "working draft" or "work in progress."

To learn the current status of any Internet-Draft, please check the `hid-abstracts.txt` listing contained in the Internet-Drafts Shadow Directories on `ftp.ietf.org`, `nic.nordu.net`, `ftp.nisc.sri.com`, or `unnari.oz.au`.

ABSTRACT

In order to support advanced interactivity as envisaged for MPEG-4 applications, this document proposes a simple RTSP-based [1] streams control framework - including necessary extensions to RTP methods syntax and semantics. Reflecting syntax and semantics of the MPEG-4 BIFS scene description [2], VRML nodes [4] and the MPEG-4 media delivery framework (DMIF) [3], in the spirit of HTMLSMIL [5], Random Access Point information, Range and Time parameters are introduced into the relevant URL(s) and related signaling methods accordingly. Two additional optional methods, R-MUTE, for Remote-MUTE, and RESUME are also proposed.

5

## 1. Motivations and Rationale

The motivations and rationale of the document is the derivation of a simple streams control framework for MPEG-4 - reflecting the syntax and semantics of the MPEG-4 scene nodes - that could possibly be introduced as extended methods syntax and semantics, in line with RTSP extension guidelines [1-section 1.5].

The attribute 'simple' alludes to the aim to provide for the minimum in support of an MPEG-4 client-to-server application equivalent to HTML/CGI or SMIL.

It is expected that in the evolution of MPEG-4 towards Java-enhanced multi-user-environments there will be in parallel other and possibly more elaborated application signaling frameworks [7].

With respect to RTSP, the main technical issues concern first a possible support for dissociation of connection management and stream control via the optional provision of a channel identifier, as an alternative to the URL in the methods. The second main issue concerns the provision of syntactical means for extended timing, random access point and range parameters. In order to support user-navigation based interactivity two additional optional methods, called R-MUTE, for Remote-MUTE, inspired from the local MUTE specified in PREMO [6], and RESUME are also proposed.

With respect to MPEG-4 Systems and DMIF the whole proposal will be fed into the ongoing Version 2 procedure. It should also be mentioned that this proposal is orthogonal to [8] which tries to position RTSP-based signaling into the DMIF environment.

## 2. Preliminary Remarks

MPEG-4 'BIFS'-scene description framework is inspired from VRML. In VRML, application specific procedural logic and state management can be implemented via Script nodes - which will be - together with Prototypes - provided only in MPEG-4 Version 2

An MPEG-4 client-server application scenario - of the type this draft is aiming at - can be characterized as follows: from the terminal side and driven by events from user interaction with the scene, a generic MPEG-4 Browser via Application Signaling requires from the Server MPEG-4 compliant streams of scene descriptions - constituting the very application - and their companions e.g. Audio-Video streams.

MPEG-4 Version 2 will introduce an advanced Interactivity Model (MPEG-J). This should lead to application specific procedural code at the terminal side, allowing e.g. for the local construction/encoding(/decoding) of BIFS updates. As the script code probably would be read as part of the scene description, the Browser could (probably) remain generic, i.e. independent of any specific application.

The signaling syntax and semantics, as discussed below, is independent of the signaling supporting mechanisms, i.e. of the procedural logic mechanisms. These mechanisms are out of the scope of this document.

### 3. Relating Application Signaling to MPEG-4 Scene Description

In VRML/MPEG-4, syntax and semantics of the nodes of a scene determine and confine the characteristics of interactivity possible. This is true for the parameters available both in 'media-playing nodes' such as MovieTexture, and in 'structure related' nodes such as Inline.

Even with future Proto and Script nodes in MPEG-4 Version 2, the expressivity of signaling, e.g. with respect to media streams, will be confined by such of the corresponding nodes in the scene.

Hence, in MPEG-4, as shortly indicated in the introduction of this document, all interactivity and in turn all application signaling has to be constructed in accordance with syntax and semantics of the relevant nodes.

#### 3.1. Media Content Playing Nodes

An object may be completely described within the scene description (BIFS) information, or may also require elementary stream data from one or more audio - visual objects, via the 'media content playing' nodes. Therefore, interactivity and corresponding signaling with respect to media objects has to be derived from the 'media content playing' nodes such as VideoObject2D, MovieTexture, etc. An application signaling method in that context would typically carry a PLAY, a PAUSE or a TEARDOWN method.

#### 3.2. Structure related nodes

Interactivity and application signaling concerning the structure of the scene, e.g. changes of a scene, will be derived from 'structure related' nodes such as Inline2D and Inline.

#### 3.3. Usage of URL's

URL as parameter of type MFString, indicating the location of the media

stream, or including a reference to an ObjectDescriptor (OD). An OD is a level of indirection that can point either to another object descriptor or to ES\_descriptors that in turn provide the references to locations of raw elementary streams associated with the node, via URL fields defined as a string of 8 bits characters (type bit(8)). Interpreted by the Browser these URL(s) will lead to the issuing of the signaling commands.

### 3.4. MPEG-4 Timing Model

A point in time at which an event occurs (change of a parameter value, defining the start or stop of a media stream, etc.) is identified by the SFTIME fields of the media content playing nodes. The SFTIME fields indicate in general a time relative to the BIFS time base that applies to the BIFS Elementary Stream that has conveyed the scene description.

The format of the SFTIME field is 64-bit double-precision floating point numbers (in ISO C floating point format) indicating a duration in seconds with respect to a reference point in time. This corresponds to an NPT - Normal Play Time - in RTSP terms, except that, here, the reference point in time (beginning of the presentation) is not expressed in GMT time but provided by the StartCompositionTimeStamp of the scene description stream.

SFTIME fields of some nodes may require 'absolute' time values, given by a "wall clock" time. The relation of the BIFS time base ticks, i.e. CTS - composition time stamp - of the BIFS Access Unit that has conveyed the respective scene description (BIFS) node, to the wall clock can be resolved, if the wall clock time is known from the receiver. This is achieved by an optional wallClockTimeStamp.

### 4. Application Signaling and the MPEG-4 media delivery framework DMIF

An MPEG4 application identifies a particular elementary stream through its Elementary Stream Id (ESid), scoped by the service session it belongs to.

When using DMIF (MPEG-4 Delivery Multimedia Framework), a 1-to-1 correspondence between each ESid and a channelHandle (chId) is realized by the DMIF layer. The stream identified by the ESid is further referred to through its channelHandle.

Dissociating connection management from stream control implies methods syntactic extensions, namely possible stream identification by a different syntactical mean other than the URL (see section 6.3). The URL will be used only for the connection management.



## 5. Extended RTSP methods syntax and semantics

### 5.1. NPT extension

The RTSP NPT format, consisting of a decimal fraction expressed in either seconds or hours, minutes, and seconds, can then be used,

```
npt-time      = "now" | npt-sec | npt-hhmmss
npt-sec       = 1*DIGIT [ "." *DIGIT ]
npt-hhmmss    = npt-hh ":" npt-mm ":" npt-ss [ "." *DIGIT ]
npt-hh        = 1*DIGIT      ; any positive number
npt-mm        = 1*2DIGIT     ; 0-59
npt-ss        = 1*2DIGIT     ; 0-59
```

However, it is necessary to provide the possibility for having the reference point in time set to the value of the startCompositionTimeStamp of the corresponding BIFS scene description stream instead of 0.0 seconds. This would mean that the beginning of the presentation is at time startCompositionTimeStamp of the corresponding BIFS scene description stream. Hence, the NPT syntax can be complemented by an optional field the wall clock time base. If this field is not present, then the default value for the reference point in time is 0.0 seconds.

```
npt-ref       = npt-hh ":" npt-mm ":" npt-ss [ "." *DIGIT ]
npt-hh        = 1*DIGIT      ; any positive number
npt-mm        = 1*2DIGIT     ; 0-59
npt-ss        = 1*2DIGIT     ; 0-59
```

### 5.2. Random Access Point (RAP) and Range extensions

Method and RAP or range information could be stored as parameters in the MFUrl class defined below in a preliminary syntax:

```
class MFUrl
{
  if (isMethod)
    SFString GlobalMethod=method;
  else
  {
    if (isOD)
    {
      bit(10) ODId;
      MFString ESId=esid ESMethod=method ESR1=ES_RAP_info1 ESR2=ES_RAP_info2;
    }
    else
    {
      SFString urlValue ODId=odid;
      MFString ESId=esid ESMethod=method ESR1=ES_RAP_info1 ESR2=ES_RAP_info2;
    }
  }
}
```

The GlobalMethod is introduced to allow for dealing with streams of all BIFS/nodes belonging to the group or whole scene.

Note that the MPEG-4 system does not, so far, provide semantical means for random access point and range information. The above class is a proposal that the authors are submitting in parallel to MPEG-4.

It is proposed here to complement the relative time and range syntax in the RTSP methods by possibly other range specifiers - also including the case of degenerated ranges specifying just a single Random Access Point. Except for the degenerated NodeID case, the 'other' ranges are still under consideration. In any case, the syntax would be:

```
other_range      = other_RAP_info#1 - other_RAP_info#2
other_rap_info   = NodeId | ...
range-specifier  = npt-range | other_range
```

### 5.3. Methods Extended Syntax

In addition to Random Access Point, Range and Time parameters, and in order to allow for dissociation of connection management and stream control, an additional syntactical mean - other than the URL mechanism - for identifying a stream must be supported.

This syntactical mechanism could be:

```
Loop(ch_identifier)
e.g. PLAY Loop(23,56,32)
```

#### 5.4. Additional Optional Methods

The two additional optional methods proposed here find a strong interest in environments with interactivity triggered by user navigation in the presented scenes (e.g. Virtual Reality in the VRML / MPEG-4 spirit). In a scene with several synchronized audio-visual streams, moving away from one audio-visual stream could allow to suspend the delivery of the stream, and coming back closer to it could resume the delivery of the stream, at a point that would be synchronized with all the other streams that have been maintained in the scene.

##### 5.4.1. R-MUTE (Remote-Mute)

The R-MUTE method is inspired from the MUTE method specified by PREMO [6]. However, in PREMO, the MUTE command suspends the presentation of the streams on the terminal but does not suspend the delivery of the streams.

The R-MUTE method would cause the stream delivery to be suspended temporarily, but a 'local' progression - on the server side - on the streams, with maintained synchronization actions, occurs without delivery of the streams. The server will hence maintain the current reading points of the on-going streams, and will then be able to resume the delivery at the corresponding random access point, when triggered by the RESUME method.

##### 5.4.2. RESUME

The RESUME method causes the re-start of the delivery of the stream, that has been previously suspended by the R-MUTE method. The delivery will be resumed at the random access point given by the server state machine, which is also dependent of the stream time base and of the time interval between the R-MUTE and the RESUME commands.

Remark: The above functionalities supposes the provision in the scene description syntax and semantic of mechanisms for routing the whole semantic of the user navigation triggered action to the media content playing nodes.

6. . Authors Addresses

Paul Christ  
Computer Center - RUS  
University of Stuttgart  
Allmandring 30  
D70550 Stuttgart, Germany.  
email: Paul.Christ@rus.uni-stuttgart.de

Christine Guillemot  
INRIA  
Campus Universitaire de Beaulieu  
35042 RENNES Cedex, FRANCE  
email: Christine.Guillemot@irisa.fr

Stefan Wesner  
Computer Center - RUS  
University of Stuttgart  
Allmandring 3a  
D70550 Stuttgart, Germany.  
email: Stefan.Wesner@rus.uni-stuttgart.de

## 7. References

- [1] H. Schulzrinne, A. Rao, R. Lanphier, 'RTSP: Real Time Streaming Protocol', RFC 2326, April 1998.
- [2] Information Technology - Coding of Audiovisual Objects - Part 1: Systems, ISO/IEC FCD 14496-1 [DRAFT], May-15, 1998.
- [3] Information technology - Generic coding of moving pictures and associated audio information - Part 6: Delivery Multimedia Integration Framework', ISO/IEC 14496-6, May-15, 1998.
- [4] VRML 97: The Virtual Reality Modeling Language, International Standard ISO/IEC 14772-1:1997
- [5] Synchronized Multimedia Integration Language (SMIL) 1.0 Specification, W3C proposed recommendation, April-9, 1998.
- [6] Information Processing Systems - Computer Graphics and Image Processing - Presentation Environments for Multimedia Objects (PREMO), Part 3: Multimedia Systems Services, ISO/IEC 14478-3.
- [7] ISO/IEC/JTC1/SC29/WG11: w/N2359 subpart 2 Verification Model of Advanced BIFS (Systems VM subpart 2) . July 98
- [8] ISO/IEC/JTC1/SC29/WG11: MPEG98/M4102, October 1998; containing: draft-balabanian-rtsp-mpeg4-dmif-00.txt . Sept. 22, 1998: Balabanian: The Role of DMIF with RTSP and MPEG-4

Internet Engineering Task Force  
 Internet Draft  
 ietf-mmusic-rtsp-02.txt  
 March 27, 1997  
 Expires: September 26, 1997

MMUSIC WG  
 H. Schulzrinne, A. Rao, R. Lanphier  
 Columbia U./Netscape/Progressive Networks

## Real Time Streaming Protocol (RTSP)

### STATUS OF THIS MEMO

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as ``work in progress''.

To learn the current status of any Internet-Draft, please check the ``lid-abstracts.txt'' listing contained in the Internet-Drafts Shadow Directories on ftp.is.co.za (Africa), nic.nordu.net (Europe), munnari.oz.au (Pacific Rim), ds.internic.net (US East Coast), or ftp.isi.edu (US West Coast).

Distribution of this document is unlimited.

### ABSTRACT

The Real Time Streaming Protocol, or RTSP, is an application-level protocol for control over the delivery of data with real-time properties. RTSP provides an extensible framework to enable controlled, on-demand delivery of real-time data, such as audio and video. Sources of data can include both live data feeds and stored clips. This protocol is intended to control multiple data delivery sessions, provide a means for choosing delivery channels such as UDP, multicast UDP and TCP, and delivery mechanisms based upon RTP (RFC 1889).

## 1 Introduction

### 1.1 Purpose

The Real-Time Streaming Protocol (RTSP) establishes and controls

either a single or several time-synchronized streams of continuous media such as audio and video. It does not typically deliver the continuous streams itself, although interleaving of the continuous media stream with the control stream is possible (see Section 9.11). In other words, RTSP acts as a "network remote control" for multimedia servers.

The set of streams to be controlled is defined by a presentation description. This memorandum does not define a format for a presentation description.

There is no notion of an RTSP connection; instead, a server maintains a session labeled by an identifier. An RTSP session is in no way tied to a transport-level connection such as a TCP connection. During an RTSP session, an RTSP client may open and close many reliable transport connections to the server to issue RTSP requests. Alternatively, it may use a connectionless transport protocol such as UDP.

The streams controlled by RTSP may use RTP [1], but the operation of RTSP does not depend on the transport mechanism used to carry continuous media.

The protocol is intentionally similar in syntax and operation to HTTP/1.1, so that extension mechanisms to HTTP can in most cases also be added to RTSP. However, RTSP differs in a number of important aspects from HTTP:

- o RTSP introduces a number of new methods and has a different protocol identifier.
- o An RTSP server needs to maintain state by default in almost all cases, as opposed to the stateless nature of HTTP. (RTSP servers and clients MAY use the HTTP state maintenance mechanism [2].)
- o Both an RTSP server and client can issue requests.
- o Data is carried out-of-band, by a different protocol. (There is an exception to this.)
- o RTSP is defined to use ISO 10646 (UTF-8) rather than ISO 8859-1, consistent with current HTML internationalization efforts [3].
- o The Request-URI always contains the absolute URI. Because of backward compatibility with a historical blunder, HTTP/1.1 carries only the absolute path in the request

This makes virtual hosting easier. However, this is incompatible with HTTP/1.1, which may be a bad idea.

The protocol supports the following operations:

Retrieval of media from media server: The client can request a presentation description via HTTP or some other method. If the presentation is being multicast, the presentation description contains the multicast addresses and ports to be used for the continuous media. If the presentation is to be sent only to the client via unicast, the client provides the destination for security reasons.

Invitation of a media server to a conference: A media server can be "invited" to join an existing conference, either to play back media into the presentation or to record all or a subset of the media in a presentation. This mode is useful for distributed teaching applications. Several parties in the conference may take turns "pushing the remote control buttons".

Addition of media to an existing presentation: Particularly for live presentations, it is useful if the server can tell the client about additional media becoming available.

RTSP requests may be handled by proxies, tunnels and caches as in HTTP/1.1.

## 1.2 Requirements

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC xxxx [4].

## 1.3 Terminology

Some of the terminology has been adopted from HTTP/1.1 [5]. Terms not listed here are defined as in HTTP/1.1.

Conference: a multiparty, multimedia presentation, where "multi" implies greater than or equal to one.

Client: The client requests continuous media data from the media server.

Connection: A transport layer virtual circuit established between two programs for the purpose of communication.

Continuous media: Data where there is a timing relationship between



source and sink, that is, the sink must reproduce the timing relationship that existed at the source. The most common examples of continuous media are audio and motion video. Continuous media can be realtime (interactive) , where there is a "tight" timing relationship between source and sink, or streaming (playback) , where the relationship is less strict.

**Participant:** Participants are members of conferences. A participant may be a machine, e.g., a media record or playback server.

**Media server:** The network entity providing playback or recording services for one or more media streams. Different media streams within a presentation may originate from different media servers. A media server may reside on the same or a different host as the web server the presentation is invoked from.

**Media parameter:** Parameter specific to a media type that may be changed while the stream is being played or prior to it.

**(Media) stream:** A single media instance, e.g., an audio stream or a video stream as well as a single whiteboard or shared application group. When using RTP, a stream consists of all RTP and RTCP packets created by a source within an RTP session. This is equivalent to the definition of a DSM-CC stream.

**Message:** The basic unit of RTSP communication, consisting of a structured sequence of octets matching the syntax defined in Section 14 and transmitted via a connection or a connectionless protocol.

**Presentation:** A set of one or more streams which the server allows the client to manipulate together. A presentation has a single time axis for all streams belonging to it. Presentations are defined by presentation descriptions (see below). A presentation description contains RTSP URIs that define which streams can be controlled individually and an RTSP URI to control the whole presentation. A movie or live concert consisting of one or more audio and video streams is an example of a presentation.

**Presentation description:** A presentation description contains information about one or more media streams within a presentation, such as the set of encodings, network addresses and information about the content. Other IETF protocols such as SDP [6] use the term "session" for a live presentation. The presentation description may take several different formats, including but not limited to the session description format SDP.

**Response:** An RTSP response. If an HTTP response is meant, that is

indicated explicitly.

**Request:** An RTSP request. If an HTTP request is meant, that is indicated explicitly.

**RTSP session:** A complete RTSP "transaction", e.g., the viewing of a movie. A session typically consist of a client setting up a transport mechanism for the continuous media stream ( SETUP), starting the stream with PLAY or RECORD and closing the stream with TEARDOWN.

#### 1.4 Protocol Properties

RTSP has the following properties:

**Extendable:** New methods and parameters can be easily added to RTSP.

**Easy to parse:** RTSP can be parsed by standard HTTP or MIME parsers.

**Secure:** RTSP re-uses web security mechanisms, either at the transport level (TLS [7]) or within the protocol itself. All HTTP authentication mechanisms such as basic [5] and digest authentication [8] are directly applicable.

**Transport-independent:** RTSP may use either an unreliable datagram protocol (UDP) [9], a reliable datagram protocol (RDP, not widely used [10]) or a reliable stream protocol such as TCP [11] as it implements application-level reliability.

**Multi-server capable:** Each media stream within a presentation can reside on a different server. The client automatically establishes several concurrent control sessions with the different media servers. Media synchronization is performed at the transport level.

**Control of recording devices:** The protocol can control both recording and playback devices, as well as devices that can alternate between the two modes ("VCR").

**Separation of stream control and conference initiation:** Stream control is divorced from inviting a media server to a conference. The only requirement is that the conference initiation protocol either provides or can be used to create a unique conference identifier. In particular, SIP [12] or H.323 may be used to invite a server to a conference.

**Suitable for professional applications:** RTSP supports frame-level accuracy through SMPTE time stamps to allow remote digital

editing.

**Presentation description neutral:** The protocol does not impose a particular presentation description or metafile format and can convey the type of format to be used. However, the presentation description must contain at least one RTSP URI.

**Proxy and firewall friendly:** The protocol should be readily handled by both application and transport-layer (SOCKS [13]) firewalls. A firewall may need to understand the SETUP method to open a "hole" for the UDP media stream.

**HTTP-friendly:** Where sensible, RTSP re-uses HTTP concepts, so that the existing infrastructure can be re-used. This infrastructure includes JEPI (the Joint Electronic Payment Initiative) for electronic payments and PICS (Platform for Internet Content Selection) for associating labels with content. However, RTSP does not just add methods to HTTP, since the controlling continuous media requires server state in most cases.

**Appropriate server control:** If a client can start a stream, it must be able to stop a stream. Servers should not start streaming to clients in such a way that clients cannot stop the stream.

**Transport negotiation:** The client can negotiate the transport method prior to actually needing to process a continuous media stream.

**Capability negotiation:** If basic features are disabled, there must be some clean mechanism for the client to determine which methods are not going to be implemented. This allows clients to present the appropriate user interface. For example, if seeking is not allowed, the user interface must be able to disallow moving a sliding position indicator.

An earlier requirement in RTSP' was multi-client capability. However, it was determined that a better approach was to make sure that the protocol is easily extensible to the multi-client scenario. Stream identifiers can be used by several control streams, so that "passing the remote" would be possible. The protocol would not address how several clients negotiate access; this is left to either a "social protocol" or some other floor control mechanism.

## 1.5 Extending RTSP

Since not all media servers have the same functionality, media

servers by necessity will support different sets of requests. For example:

- o A server may only be capable of playback, not recording and thus has no need to support the RECORD request.
- o A server may not be capable of seeking (absolute positioning), say, if it is to support live events only.
- o Some servers may not support setting stream parameters and thus not support GET\_PARAMETER and SET\_PARAMETER.

A server SHOULD implement all header fields described in Section 11.

It is up to the creators of presentation descriptions not to ask the impossible of a server. This situation is similar in HTTP/1.1, where the methods described in [H19.6] are not likely to be supported across all servers.

RTSP can be extended in three ways, listed in order of the magnitude of changes supported:

- o Existing methods can be extended with new parameters, as long as these parameters can be safely ignored by the recipient. (This is equivalent to adding new parameters to an HTML tag.)
- o New methods can be added. If the recipient of the message does not understand the request, it responds with error code 501 (Not implemented) and the sender can then attempt an earlier, less functional version.
- o A new version of the protocol can be defined, allowing almost all aspects (except the position of the protocol version number) to change.

## 1.6 Overall Operation

Each presentation and media stream may be identified by an RTSP URL. The overall presentation and the properties of the media the presentation is made up of are defined by a presentation description file, the format of which is outside the scope of this specification. The presentation description file may be obtained by the client using HTTP or other means such as email and may not necessarily be stored on the media server.

For the purposes of this specification, a presentation description is assumed to describe one or more presentations, each of which maintains a common time axis. For simplicity of exposition and

without loss of generality, it is assumed that the presentation description contains exactly one such presentation. A presentation may contain several media streams.

The presentation description file contains a description of the media streams making up the presentation, including their encodings, language, and other parameters that enable the client to choose the most appropriate combination of media. In this presentation description, each media stream that is individually controllable by RTSP is identified by an RTSP URL, which points to the media server handling that particular media stream and names the stream stored on that server. Several media streams can be located on different servers; for example, audio and video streams can be split across servers for load sharing. The description also enumerates which transport methods the server is capable of.

Besides the media parameters, the network destination address and port need to be determined. Several modes of operation can be distinguished:

Unicast: The media is transmitted to the source of the RTSP request, with the port number chosen by the client. Alternatively, the media is transmitted on the same reliable stream as RTSP.

Multicast, server chooses address: The media server picks the multicast address and port. This is the typical case for a live or near-media-on-demand transmission.

Multicast, client chooses address: If the server is to participate in an existing multicast conference, the multicast address, port and encryption key are given by the conference description, established by means outside the scope of this specification.

## 1.7 RTSP States

RTSP controls a stream which may be sent via a separate protocol, independent of the control channel. For example, RTSP control may occur on a TCP connection while the data flows via UDP. Thus, data delivery continues even if no RTSP requests are received by the media server. Also, during its lifetime, a single media stream may be controlled by RTSP requests issued sequentially on different TCP connections. Therefore, the server needs to maintain "session state" to be able to correlate RTSP requests with a stream. The state transitions are described in Section A.

Many methods in RTSP do not contribute to state. However, the following play a central role in defining the allocation and usage of stream resources on the server: SETUP, PLAY, RECORD, PAUSE, and

TEARDOWN.

SETUP: Causes the server to allocate resources for a stream and start an RTSP session.

PLAY and RECORD: Starts data transmission on a stream allocated via SETUP.

PAUSE: Temporarily halts a stream, without freeing server resources.

TEARDOWN: Frees resources associated with the stream. The RTSP session ceases to exist on the server.

## 1.8 Relationship with Other Protocols

RTSP has some overlap in functionality with HTTP. It also may interact with HTTP in that the initial contact with streaming content is often to be made through a web page. The current protocol specification aims to allow different hand-off points between a web server and the media server implementing RTSP. For example, the presentation description can be retrieved using HTTP or RTSP. Having the presentation description be returned by the web server makes it possible to have the web server take care of authentication and billing, by handing out a presentation description whose media identifier includes an encrypted version of the requestor's IP address and a timestamp, with a shared secret between web and media server.

However, RTSP differs fundamentally from HTTP in that data delivery takes place out-of-band, in a different protocol. HTTP is an asymmetric protocol, where the client issues requests and the server responds. In RTSP, both the media client and media server can issue requests. RTSP requests are also not stateless, in that they may set parameters and continue to control a media stream long after the request has been acknowledged.

Re-using HTTP functionality has advantages in at least two areas, namely security and proxies. The requirements are very similar, so having the ability to adopt HTTP work on caches, proxies and authentication is valuable.

While most real-time media will use RTP as a transport protocol, RTSP is not tied to RTP.

RTSP assumes the existence of a presentation description format that can express both static and temporal properties of a presentation containing several media streams.

## 2 Notational Conventions

Since many of the definitions and syntax are identical to HTTP/1.1, this specification only points to the section where they are defined rather than copying it. For brevity, [HX.Y] is to be taken to refer to Section X.Y of the current HTTP/1.1 specification (RFC 2068).

All the mechanisms specified in this document are described in both prose and an augmented Backus-Naur form (BNF) similar to that used in RFC 2068 [H2.1]. It is described in detail in [14].

In this draft, we use indented and smaller-type paragraphs to provide background and motivation. Some of these paragraphs are marked with HS, AR and RL, designating opinions and comments by the individual authors which may not be shared by the co-authors and require resolution.

## 3 Protocol Parameters

### 3.1 RTSP Version

applies, with HTTP replaced by RTSP.

### 3.2 RTSP URL

The "rtsp" and "rtspu" schemes are used to refer to network resources via the RTSP protocol. This section defines the scheme-specific syntax and semantics for RTSP URLs.

```
rtsp_URL = ( "rtsp:" | "rtspu:" ) "://" host [ ":" port ] [abs_path]
host      = <A legal Internet host domain name or IP address
           (in dotted decimal form), as defined by Section 2.1
           of RFC 1123>
port      = *DIGIT
```

abs\_path is defined in [H3.2.1].

Note that fragment and query identifiers do not have a well-defined meaning at this time, with the interpretation left to the RTSP server.

The scheme rtsp requires that commands are issued via a reliable protocol (within the Internet, TCP), while the scheme rtspu identifies an unreliable protocol (within the Internet, UDP).

If the port is empty or not given, port 554 is assumed. The semantics are that the identified resource can be controlled by RTSP at the server listening for TCP (scheme "rtsp") connections or UDP (scheme "rtspu") packets on that port of host, and the Request-URI for the resource is rtsp\_URL.

The use of IP addresses in URLs SHOULD be avoided whenever possible (see RFC 1924 [15]).

A presentation or a stream is identified by a textual media identifier, using the character set and escape conventions [H3.2] of URLs [16]. Requests described in Section 9 can refer to either the whole presentation or an individual stream within the presentation. Note that some methods can only be applied to streams, not presentations and vice versa. A specific instance of a presentation or stream, e.g., one of several concurrent transmissions of the same content, an RTSP session, is indicated by the Session header field (Section 11.26) where needed.

For example, the RTSP URL

```
rtsp://media.example.com:554/twister/audiotrack
```

identifies the audio stream within the presentation "twister", which can be controlled via RTSP requests issued over a TCP connection to port 554 of host media.example.com

This does not imply a standard way to reference streams in URLs. The presentation description defines the hierarchical relationships in the presentation and the URLs for the individual streams. A presentation description may name a stream 'a.mov' and the whole presentation 'b.mov'.

The path components of the RTSP URL are opaque to the client and do not imply any particular file system structure for the server.

This decoupling also allows presentation descriptions to be used with non-RTSP media control protocols, simply by replacing the scheme in the URL.

### 3.3 Conference Identifiers

Conference identifiers are opaque to RTSP and are encoded using standard URI encoding methods (i.e., LWS is escaped with %). They can contain any octet value. The conference identifier MUST be globally



unique. For H.323, the conferenceID value is to be used.

conference-id = 1\*OCTET ; LWS must be URL-escaped

Conference identifiers are used to allow to allow RTSP sessions to obtain parameters from multimedia conferences the media server is participating in. These conferences are created by protocols outside the scope of this specification, e.g., H.323 [17] or SIP [12]. Instead of the RTSP client explicitly providing transport information, for example, it asks the media server to use the values in the conference description instead. If the conference participant inviting the media server would only supply a conference identifier which is unique for that inviting party, the media server could add an internal identifier for that party, e.g., its Internet address. However, this would prevent that the conference participant and the initiator of the RTSP commands are two different entities.

### 3.4 SMPTE Relative Timestamps

A SMPTE relative time-stamp expresses time relative to the start of the clip. Relative timestamps are expressed as SMPTE time codes for frame-level access accuracy. The time code has the format  
hours:minutes:seconds.frames

with the origin at the start of the clip. For NTSC, the frame rate is 29.97 frames per second. This is handled by dropping the first frame index of every minute, except every tenth minute. If the frame value is zero, it may be omitted.

smpte-range = "smpte" "=" smpte-time "-" [ smpte-time ]  
smpte-time = 1\*2DIGIT ":" 1\*2DIGIT ":" 1\*2DIGIT [ "." 1\*2DIGIT ]

#### Examples:

smpte=10:12:33.40-  
smpte=10:7:33-  
smpte=10:7:0-10:7:33

### 3.5 Normal Play Time

Normal play time (NPT) indicates the stream absolute position relative to the beginning of the presentation, measured in seconds and microseconds. The beginning of a presentation corresponds to 0 seconds and 0 microseconds. Negative values are not defined. The microsecond field is always less than 1,000,000. NPT is defined as in DSM-CC: "Intuitively, NPT is the clock the viewer associates with a program. It is often digitally displayed on a VCR. NPT advances normally when in normal play mode (scale = 1), advances at a faster rate when in fast scan forward (high positive scale ratio), decrements when in scan reverse (high negative scale ratio) and is fixed in pause mode. NPT is [logically] equivalent to SMPTE time codes." [18]

```
npt-range = "npt" "=" npt-time "-" [ npt-time ]
npt-time  = 1*DIGIT [ ":" *DIGIT ]
```

#### Examples:

```
npt=123:45-125
```

### 3.6 Absolute Time

Absolute time is expressed as ISO 8601 timestamps, using UTC (GMT). Fractions of a second may be indicated.

```
utc-range = "clock" "=" utc-time "-" [ utc-time ]
utc-time  = utc-date "T" utc-time "Z"
utc-date  = 8DIGIT ; < YYYYMMDD >
utc-time  = 6DIGIT [ "." fraction ] ; < HHMMSS.fraction >
```

Example for November 8, 1996 at 14h37 and 20 and a quarter seconds UTC:

```
19961108T143720.25Z
```

#### Example

## 4 RTSP Message

RTSP is a text-based protocol and uses the ISO 10646 character set in UTF-8 encoding (RFC 2044). Lines are terminated by CRLF, but receivers should be prepared to also interpret CR and LF by themselves as line terminators.

Text-based protocols make it easier to add optional parameters in a self-describing manner. Since the number of parameters and the frequency of commands is low, processing efficiency is not a concern. Text-based protocols, if done carefully, also allow easy implementation of research prototypes in scripting languages such as Tcl, Visual Basic and Perl.

The 10646 character set avoids tricky character set switching, but is invisible to the application as long as US-ASCII is being used. This is also the encoding used for RTP. ISO 8859-1 translates directly into Unicode, with a high-order octet of zero. ISO 8859-1 characters with the most-significant bit set are represented as 1100001x10xxxxxx.

RTSP messages can be carried over any lower-layer transport protocol that is 8-bit clean.

Requests contain methods, the object the method is operating upon and parameters to further describe the method. Methods are idempotent, unless otherwise noted. Methods are also designed to require little or no state maintenance at the media server.

#### 4.1 Message Types

See [H4.1]

#### 4.2 Message Headers

See [H4.2]

#### 4.3 Message Body

See [H4.3]

#### 4.4 Message Length

When a message-body is included with a message, the length of that body is determined by one of the following (in order of precedence):

1. Any response message which MUST NOT include a message-body (such as the 1xx, 204, and 304 responses) is always

terminated by the first empty line after the header fields, regardless of the entity-header fields present in the message. (Note: An empty line consists of only CRLF.)

2. If a Content-Length header field (section 11.12) is present, its value in bytes represents the length of the message-body. If this header field is not present, a value of zero is assumed.
3. By the server closing the connection. (Closing the connection cannot be used to indicate the end of a request body, since that would leave no possibility for the server to send back a response.)

Note that RTSP does not (at present) support the HTTP/1.1 "chunked" transfer coding and requires the presence of the Content-Length header field.

Given the moderate length of presentation descriptions returned, the server should always be able to determine its length, even if it is generated dynamically, making the chunked transfer encoding unnecessary. Even though Content-Length must be present if there is any entity body, the rules ensure reasonable behavior even if the length is not given explicitly.

## 5 Request

A request message from a client to a server or vice versa includes, within the first line of that message, the method to be applied to the resource, the identifier of the resource, and the protocol version in use.

```
Request = Request-line CRLF
        *request-header
        CRLF
        [ message-body ]
```

```
Request-Line = Method SP Request-URI SP RTSP-Version SP seq-no CRLF
```

```
Method = "DESCRIBE"           ; Section
        | "GET_PARAMETER"     ; Section
        | "OPTIONS"           ; Section
        | "PAUSE"              ; Section
        | "PLAY"               ; Section
        | "RECORD"             ; Section
```

```

| "REDIRECT"           ; Section
| "SETUP"              ; Section
| "SET_PARAMETER"      ; Section
| "TEARDOWN"           ; Section
| extension-method

```

extension-method = token

Request-URI = "\*" | absolute\_URI

RTSP-Version = "RTSP" "/" 1\*DIGIT "." 1\*DIGIT

seq-no = 1\*DIGIT

Note that in contrast to HTTP/1.1, RTSP requests always contain the absolute URL (that is, including the scheme, host and port) rather than just the absolute path.

The asterisk "\*" in the Request-URI means that the request does not apply to a particular resource, but to the server itself, and is only allowed when the method used does not necessarily apply to a resource. One example would be

OPTIONS \* RTSP/1.0

## 6 Response

[H6] applies except that HTTP-Version is replaced by RTSP-Version define some HTTP codes. The valid response codes and the methods they can be used with are defined in the table 1.

After receiving and interpreting a request message, the recipient responds with an RTSP response message.

```

Response = Status-Line           ; Section
          *( general-header      ; Section
            | response-header    ; Section
            | entity-header )    ; Section
          CRLF
          [ message-body ]       ; Section

```

## 6.1 Status-Line

The first line of a Response message is the Status-Line , consisting of the protocol version followed by a numeric status code, the sequence number of the corresponding request and the textual phrase associated with the status code, with each element separated by SP characters. No CR or LF is allowed except in the final CRLF sequence. Note that the addition of a

Status-Line = RTSP-Version SP Status-Code SP seq-no SP Reason-Phrase CRLF

### 6.1.1 Status Code and Reason Phrase

The Status-Code element is a 3-digit integer result code of the attempt to understand and satisfy the request. These codes are fully defined in section10. The Reason-Phrase is intended to give a short textual description of the Status-Code. The Status-Code is intended for use by automata and the Reason-Phrase is intended for the human user. The client is not required to examine or display the Reason-Phrase

The first digit of the Status-Code defines the class of response. The last two digits do not have any categorization role. There are 5 values for the first digit:

- o 1xx: Informational - Request received, continuing process
- o 2xx: Success - The action was successfully received, understood, and accepted
- o 3xx: Redirection - Further action must be taken in order to complete the request
- o 4xx: Client Error - The request contains bad syntax or cannot be fulfilled
- o 5xx: Server Error - The server failed to fulfill an apparently valid request

The individual values of the numeric status codes defined for RTSP/1.0, and an example set of corresponding Reason-Phrase below. The reason phrases listed here are only recommended -- they may be replaced by local equivalents without affecting the protocol. Note that RTSP adopts most HTTP/1.1 status codes and adds RTSP-specific status codes in the starting at 450 to avoid conflicts with newly

defined HTTP status codes.

```
Status-Code    = "100"    ; Continue
                  | "200"    ; OK
                  | "201"    ; Created
                  | "300"    ; Multiple Choices
                  | "301"    ; Moved Permanently
                  | "302"    ; Moved Temporarily
                  | "303"    ; See Other
                  | "304"    ; Not Modified
                  | "305"    ; Use Proxy
                  | "400"    ; Bad Request
                  | "401"    ; Unauthorized
                  | "402"    ; Payment Required
                  | "403"    ; Forbidden
                  | "404"    ; Not Found
                  | "405"    ; Method Not Allowed
                  | "406"    ; Not Acceptable
                  | "407"    ; Proxy Authentication Required
                  | "408"    ; Request Time-out
                  | "409"    ; Conflict
                  | "410"    ; Gone
                  | "411"    ; Length Required
                  | "412"    ; Precondition Failed
                  | "413"    ; Request Entity Too Large
                  | "414"    ; Request-URI Too Large
                  | "415"    ; Unsupported Media Type
                  | "451"    ; Parameter Not Understood
                  | "452"    ; Conference Not Found
                  | "453"    ; Not Enough Bandwidth
                  | "45x"    ; Session Not Found
                  | "45x"    ; Method Not Valid in This State
                  | "45x"    ; Header Field Not Valid for Resource
                  | "45x"    ; Invalid Range
                  | "45x"    ; Parameter Is Read-Only
                  | "500"    ; Internal Server Error
                  | "501"    ; Not Implemented
                  | "502"    ; Bad Gateway
                  | "503"    ; Service Unavailable
                  | "504"    ; Gateway Time-out
                  | "505"    ; HTTP Version not supported
                  | extension-code
```

extension-code = 3DIGIT

Reason-Phrase = \*<TEXT, excluding CR, LF>

RTSP status codes are extensible. RTSP applications are not required to understand the meaning of all registered status codes, though such understanding is obviously desirable. However, applications **MUST** understand the class of any status code, as indicated by the first digit, and treat any unrecognized response as being equivalent to the x00 status code of that class, with the exception that an unrecognized response **MUST NOT** be cached. For example, if an unrecognized status code of 431 is received by the client, it can safely assume that there was something wrong with its request and treat the response as if it had received a 400 status code. In such cases, user agents **SHOULD** present to the user the entity returned with the response, since that entity is likely to include human-readable information which will explain the unusual status.

### 6.1.2 Response Header Fields

The response-header fields allow the request recipient to pass additional information about the response which cannot be placed in the Status-Line server and about further access to the resource identified by the Request-URI

```
response-header = Location           ; Section
                  | Proxy-Authenticate ; Section
                  | Public             ; Section
                  | Retry-After        ; Section
                  | Server              ; Section
                  | Vary                ; Section
                  | WWW-Authenticate   ; Section
```

Response-header field names can be extended reliably only in combination with a change in the protocol version. However, new or experimental header fields **MAY** be given the semantics of response-header fields if all parties in the communication recognize them to be response-header fields. Unrecognized header fields are treated as entity-header fields.

## 7 Entity

Request and Response messages **MAY** transfer an entity if not otherwise restricted by the request method or response status code. An entity consists of entity-header fields and an entity-body, although some responses will only include the entity-headers.

In this section, both sender and recipient refer to either the client



Code	reason	
100	Continue	all
200	OK	all
201	Created	RECORD
300	Multiple Choices	all
301	Moved Permanently	all
302	Moved Temporarily	all
303	See Other	all
305	Use Proxy	all
400	Bad Request	all
401	Unauthorized	all
402	Payment Required	all
403	Forbidden	all
404	Not Found	all
405	Method Not Allowed	all
406	Not Acceptable	all
407	Proxy Authentication Required	all
408	Request Timeout	all
409	Conflict	
410	Gone	all
411	Length Required	SETUP
412	Precondition Failed	
413	Request Entity Too Large	SETUP
414	Request-URI Too Long	all
415	Unsupported Media Type	SETUP
45x	Only Valid for Stream	SETUP
45x	Invalid parameter	SETUP
45x	Not Enough Bandwidth	SETUP
45x	Illegal Conference Identifier	SETUP
45x	Illegal Session Identifier	PLAY, RECORD, TEARDOWN
45x	Parameter Is Read-Only	SET_PARAMETER
45x	Header Field Not Valid	all
500	Internal Server Error	all
501	Not Implemented	all
502	Bad Gateway	all
503	Service Unavailable	all
504	Gateway Timeout	all
505	RTSP Version Not Supported	all

Table 1: Status codes and their usage with RTSP methods

or the server, depending on who sends and who receives the entity.

## 7.1 Entity Header Fields

Entity-header fields define optional metainformation about the entity-body or, if no body is present, about the resource identified by the request.

```
entity-header = Allow                ; Section 14.7
                | Content-Encoding    ; Section 14.12
                | Content-Language    ; Section 14.13
                | Content-Length      ; Section 14.14
                | Content-Type        ; Section 14.18
                | Expires              ; Section 14.21
                | Last-Modified       ; Section 14.29
                | extension-header
```

extension-header = message-header

The extension-header mechanism allows additional entity-header fields to be defined without changing the protocol, but these fields cannot be assumed to be recognizable by the recipient. Unrecognized header fields SHOULD be ignored by the recipient and forwarded by proxies.

## 7.2 Entity Body

See [H7.2]

## 8 Connections

RTSP requests can be transmitted in several different ways:

- o persistent transport connections used for several request-response transactions;
- o one connection per request/response transaction;
- o connectionless mode.

The type of transport connection is defined by the RTSP URI (Section 3.2). For the scheme "rtsp", a persistent connection is assumed, while the scheme "rtspu" calls for RTSP requests to be sent without setting up a connection.

Unlike HTTP, RTSP allows the media server to send requests to the

media client. However, this is only supported for persistent connections, as the media server otherwise has no reliable way of reaching the client. Also, this is the only way that requests from media server to client are likely to traverse firewalls.

## 8.1 Pipelining

A client that supports persistent connections or connectionless mode MAY "pipeline" its requests (i.e., send multiple requests without waiting for each response). A server MUST send its responses to those requests in the same order that the requests were received.

## 8.2 Reliability and Acknowledgements

Requests are acknowledged by the receiver unless they are sent to a multicast group. If there is no acknowledgement, the sender may resend the same message after a timeout of one round-trip time (RTT). The round-trip time is estimated as in TCP (RFC TBD), with an initial round-trip value of 500 ms. An implementation MAY cache the last RTT measurement as the initial value for future connections. If a reliable transport protocol is used to carry RTSP, the timeout value MAY be set to an arbitrarily large value.

This can greatly increase responsiveness for proxies operating in local-area networks with small RTTs. The mechanism is defined such that the client implementation does not have to be aware of whether a reliable or unreliable transport protocol is being used. It is probably a bad idea to have two reliability mechanisms on top of each other, although the RTSP RTT estimate is likely to be larger than the TCP estimate.

Each request carries a sequence number, which is incremented by one for each request transmitted. If a request is repeated because of lack of acknowledgement, the sequence number is incremented.

This avoids ambiguities when computing round-trip time estimates. [TBD: An initial sequence number negotiation needs to be added for UDP; otherwise, a new stream connection may see a request be acknowledged by a delayed response from an earlier "connection". This handshake can be avoided with a sequence number containing a timestamp of sufficiently high resolution.]

The reliability mechanism described here does not protect against reordering. This may cause problems in some instances. For example, a TEARDOWN followed by a PLAY has quite a different effect than the reverse. Similarly, if a PLAY request arrives before all parameters

are set due to reordering, the media server would have to issue an error indication. Since sequence numbers for retransmissions are incremented (to allow easy RTT estimation), the receiver cannot just ignore out-of-order packets. [TBD: This problem could be fixed by including both a sequence number that stays the same for retransmissions and a timestamp for RTT estimation.]

Systems implementing RTSP MUST support carrying RTSP over TCP and MAY support UDP. The default port for the RTSP server is 554 for both UDP and TCP.

A number of RTSP packets destined for the same control end point may be packed into a single lower-layer PDU or encapsulated into a TCP stream. RTSP data MAY be interleaved with RTP and RTCP packets. Unlike HTTP, an RTSP method header MUST contain a Content-Length whenever that method contains a payload. Otherwise, an RTSP packet is terminated with an empty line immediately following the method header.

## 9 Method Definitions

The method token indicates the method to be performed on the resource identified by the Request-URI case-sensitive. New methods may be defined in the future. Method names may not start with a \$ character (decimal 24) and must be a token

method	direction	object	requirement
DESCRIBE	C -> S, S -> C	P,S	recommended
GET_PARAMETER	C -> S, S -> C	P,S	optional
OPTIONS	C -> S	P,S	required
PAUSE	C -> S	P,S	recommended
PLAY	C -> S	P,S	required
RECORD	C -> S	P,S	optional
REDIRECT	S -> C	P,S	optional
SETUP	C -> S	S	required
SET_PARAMETER	C -> S, S -> C	P,S	optional
TEARDOWN	C -> S	P,S	required

Table 2: Overview of RTSP methods, their direction, and what objects (P: presentation, S: stream) they operate on

Notes on Table 2: PAUSE is recommend, but not required in that a fully functional server can be built that does not support this method, for example, for live feeds. If a server does not support a particular method, it MUST return "501 Not Implemented" and a client

SHOULD not try this method again for this server.

## 9.1 OPTIONS

The behavior is equivalent to that described in [H9.2]. An OPTIONS request may be issued at any time, e.g., if the client is about to try a non-standard request. It does not influence server state.

In addition, if the optional Require header is present, option tags within the header indicate features needed by the requestor that are not required at the version level of the protocol.

Example 1:

```
C->S: OPTIONS * RTSP/1.0 1
      Require: implicit-play, record-feature
      Transport-Require: switch-to-udp-control, gzipped-messages
```

Note that these are fictional features (though we may want to make them real one day).

Example 2 (using RFC2069-style authentication only as an example):

```
S->C: OPTIONS * RTSP/1.0 1
      Authenticate: Digest realm="testrealm@host.com",
        nonce="dcd98b7102dd2f0e8b11d0f600bfb0c093",
        opaque="5ccc069c403ebaf9f0171e9517f40e41"

S->C: RTSP/1.0 200 1 OK
      Date: 23 Jan 1997 15:35:06 GMT
      Nack-Transport-Require: switch-to-udp-control
```

Note that these are fictional features (though we may want to make them real one day).

Example 2 (using RFC2069-style authentication only as an example):

```
C->S: RTSP/1.0 401 1 Unauthorized
      Authorization: Digest username="Mufasa",
        realm="testrealm@host.com",
        nonce="dcd98b7102dd2f0e8b11d0f600bfb0c093",
        uri="/dir/index.html",
```

```
response="e966c932a9242554e42c8ee200cec7f6",
opaque="5ccc069c403ebaf9f0171e9517f40e41"
```

## 9.2 DESCRIBE

The DESCRIBE method retrieves the description of a presentation or media object identified by the request URL from a server. It may use the Accept header to specify the description formats that the client understands. The server responds with a description of the requested resource. Alternatively, the server may "push" a new description to the client, for example, if a new stream has become available. If a new media stream is added to a presentation (e.g., during a live presentation), the whole presentation description should be sent again, rather than just the additional components, so that components can be deleted.

Example:

```
C->S: DESCRIBE rtsp://server.example.com/fizzle/foo RTSP/1.0 312
      Accept: application/sdp, application/rtsp, application/mpeg
```

```
S->C: RTSP/1.0 200 312 OK
      Date: 23 Jan 1997 15:35:06 GMT
      Content-Type: application/sdp
      Content-Length: 376
```

```
v=0
o=mhandley 2890844526 2890842807 IN IP4 126.16.64.4
s=SDP Seminar
i=A Seminar on the session description protocol
u=http://www.cs.ucl.ac.uk/staff/M.Handley/sdp.03.ps
e=mjh@isi.edu (Mark Handley)
c=IN IP4 224.2.17.12/127
t=2873397496 2873404696
a=recvonly
m=audio 3456 RTP/AVP 0
m=video 2232 RTP/AVP 31
m=whiteboard 32416 UDP WB
a=orient:portrait
```

or

```
S->C: RTSP/1.0 200 312 OK
      Date: 23 Jan 1997 15:35:06 GMT
```

Content-Type: application/rtsl  
Content-Length: 2782

<2782 octets of data containing stream description>

Server to client example:

S->C: DESCRIBE /twister RTSP/1.0 902  
Session: 1234  
Content-Type: application/rtsl  
  
new RTSL presentation description

### 9.3 SETUP

The SETUP request for a URI specifies the transport mechanism to be used for the streamed media. A client can issue a SETUP request for a stream that is already playing to change transport parameters. For the benefit of any intervening firewalls, a client must indicate the transport parameters even if it has no influence over these parameters, for example, where the server advertises a fixed multicast address.

This avoids having firewall to parse numerous different presentation description formats, for information which is irrelevant.

If the optional Require header is present, option tags within the header indicate features needed by the requestor that are not required at the version level of the protocol. The Transport-Require header is used to indicate proxy-sensitive features that MUST be stripped by the proxy to the server if not supported. Furthermore, any Transport-Require header features that are not supported by the proxy MUST be negatively acknowledged by the proxy to the client if not supported.

HS: In my opinion, the Require header should be replaced by PEP since PEP is standards-track, has more functionality and somebody already did the work.

The Transport header specifies the transport parameters acceptable to the client for data transmission; the response will contain the transport parameters selected by the server.

C->S: SETUP foo/bar/baz.rm RTSP/1.0 302  
Transport: rtp/udp;port=458

S->C: RTSP/1.0 200 302 OK  
Date: 23 Jan 1997 15:35:06 GMT  
Transport: cush/udp;port=458

#### 9.4 PLAY

The PLAY method tells the server to start sending data via the mechanism specified in SETUP. A client MUST NOT issue a PLAY request until any outstanding SETUP requests have been acknowledged as successful.

The PLAY request positions the normal play time to the beginning of the range specified and delivers stream data until the end of the range is reached. PLAY requests may be pipelined (queued); a server MUST queue PLAY requests to be executed in order. That is, a PLAY request arriving while a previous PLAY request is still active is delayed until the first has been completed.

This allows precise editing. For example, regardless of how closely spaced the two PLAY commands in the example below arrive, the server will play first second 10 through 15 and then, immediately following, seconds 20 to 25 and finally seconds 30 through the end.

C->S: PLAY rtsp://audio.example.com/audio RTSP/1.0 835  
Range: npt=10-15

C->S: PLAY rtsp://audio.example.com/audio RTSP/1.0 836  
Range: npt=20-25

C->S: PLAY rtsp://audio.example.com/audio RTSP/1.0 837  
Range: npt=30-

See the description of the PAUSE request for further examples.

A PLAY request without a Range header is legal. It starts playing a stream from the beginning unless the stream has been paused. If a stream has been paused via PAUSE, stream delivery resumes at the pause point. If a stream is playing, such a PLAY request causes no further action and can be used by the client to test server liveness.



The Range header may also contain a time parameter. This parameter specifies a time in UTC at which the playback should start. If the message is received after the specified time, playback is started immediately. The time parameter may be used to aid in synchronisation of streams obtained from different sources.

For a on-demand stream, the server replies back with the actual range that will be played back. This may differ from the requested range if alignment of the requested range to valid frame boundaries is required for the media source. If no range is specified in the request, the current position is returned in the reply. The unit of the range in the reply is the same as that in the request.

After playing the desired range, the presentation is automatically paused, as if a PAUSE request had been issued.

The following example plays the whole presentation starting at SMPTE time code 0:10:20 until the end of the clip. The playback is to start at 15:36 on 23 Jan 1997.

```
C->S: PLAY rtsp://audio.example.com/twister.en RTSP/1.0 833
      Range: smpte=0:10:20-;time=19970123T153600Z
```

```
S->C: RTSP/1.0 200 833 OK
      Date: 23 Jan 1997 15:35:06 GMT
      Range: smpte=0:10:22-;time=19970123T153600Z
```

For playing back a recording of a live presentation, it may be desirable to use clock units:

```
C->S: PLAY rtsp://audio.example.com/meeting.en RTSP/1.0 835
      Range: clock=19961108T142300Z-19961108T143520Z
```

```
S->C: RTSP/1.0 200 833 OK
      Date: 23 Jan 1997 15:35:06 GMT
```

A media server only supporting playback MUST support the smpte format and MAY support the clock format.

## 9.5 PAUSE

The PAUSE request causes the stream delivery to be interrupted (halted) temporarily. If the request URL names a stream, only playback and recording of that stream is halted. For example, for audio, this is equivalent to muting. If the request URL names a presentation or group of streams, delivery of all currently active streams within the presentation or group is halted. After resuming playback or recording, synchronization of the tracks MUST be maintained. Any server resources are kept.

The PAUSE request may contain a Range header specifying when the stream or presentation is to be halted. The header must contain exactly one value rather than a time range. The normal play time for the stream is set to that value. The pause request becomes effective the first time the server is encountering the time point specified. If this header is missing, stream delivery is interrupted immediately on receipt of the message.

For example, if the server has play requests for ranges 10 to 15 and 20 to 29 pending and then receives a pause request for NPT 21, it would start playing the second range and stop at NPT 21. If the pause request is for NPT 12 and the server is playing at NPT 13 serving the first play request, it stops immediately. If the pause request is for NPT 16, it stops after completing the first play request and discards the second play request.

As another example, if a server has received requests to play ranges 10 to 15 and then 13 to 20, that is, overlapping ranges, the PAUSE request for NPT=14 would take effect while playing the first range, with the second PLAY request effectively being ignored, assuming the PAUSE request arrives before the server has started playing the second, overlapping range. Regardless of when the PAUSE request arrives, it sets the NPT to 14.

If the server has already sent data beyond the time specified in the Range header, a PLAY would still resume at that point in time, as it is assumed that the client has discarded data after that point. This ensures continuous pause/play cycling without gaps.

Example:

```
C->S: PAUSE /fizzle/foo RTSP/1.0 834
```

```
S->C: RTSP/1.0 200 834 OK
      Date: 23 Jan 1997 15:35:06 GMT
```

## 9.6 TEARDOWN

Stop the stream delivery for the given URI, freeing the resources associated with it. If the URI is the root node for this presentation, any RTSP session identifier associated with the session is no longer valid. Unless all transport parameters are defined by the session description, a SETUP request has to be issued before the session can be played again.

Example:

```
C->S: TEARDOWN /fizzle/foo RTSP/1.0 892
```

```
S->C: RTSP/1.0 200 892 OK
```

## 9.7 GET\_PARAMETER

The requests retrieves the value of a parameter of a presentation or stream specified in the URI. Multiple parameters can be requested in the message body using the content type text/rtsp-parameters. Note that parameters include server and client statistics. IANA registers parameter names for statistics and other purposes. GET\_PARAMETER with no entity body may be used to test client or server liveness ("ping").

Example:

```
S->C: GET_PARAMETER /fizzle/foo RTSP/1.0 431
      Content-Type: text/rtsp-parameters
      Session: 1234
      Content-Length: 15
```

```
      packets_received
      jitter
```

```
C->S: RTSP/1.0 200 431 OK
      Content-Length: 46
      Content-Type: text/rtsp-parameters
```

```
      packets_received: 10
      jitter: 0.3838
```

## 9.8 SET\_PARAMETER

This method requests to set the value of a parameter for a presentation or stream specified by the URI.

A request SHOULD only contain a single parameter to allow the client to determine why a particular request failed. A server MUST allow a parameter to be set repeatedly to the same value, but it MAY disallow changing parameter values.

Note: transport parameters for the media stream MUST only be set with the SETUP command.

Restricting setting transport parameters to SETUP is for the benefit of firewalls.

The parameters are split in a fine-grained fashion so that there can be more meaningful error indications. However, it may make sense to allow the setting of several parameters if an atomic setting is desirable. Imagine device control where the client does not want the camera to pan unless it can also tilt to the right angle at the same time.

A SET\_PARAMETER request without parameters can be used as a way to detect client or server liveness.

Example:

```
C->S: SET_PARAMETER /fizzle/foo RTSP/1.0 421
```

```
Content-type: text/rtsp-parameters
```

```
fooparam: foostuff
```

```
barparam: barstuff
```

```
S->C: RTSP/1.0 450 421 Invalid Parameter
```

```
Content-Length: 6
```

```
barparam
```

## 9.9 REDIRECT

A redirect request informs the client that it must connect to another server location. It contains the mandatory header Location, which indicates that the client should issue a DESCRIBE for that URL. It

may contain the parameter `Range`, which indicates when the redirection takes effect.

This example request redirects traffic for this URI to the new server at the given play time:

```
S->C: REDIRECT /fizzle/foo RTSP/1.0 732
      Location: rtsp://bigserver.com:8001
      Range: clock=19960213T143205Z-
```

## 9.10 RECORD

This method initiates recording a range of media data according to the presentation description. The timestamp reflects start and end time (UTC). If no time range is given, use the start or end time provided in the presentation description. If the session has already started, commence recording immediately. The `Conference` header is mandatory.

The server decides whether to store the recorded data under the request-URI or another URI. If the server does not use the request-URI, the response **SHOULD** be 201 (Created) and contain an entity which describes the status of the request and refers to the new resource, and a `Location` header.

A media server supporting recording of live presentations **MUST** support the clock range format; the `smpte` format does not make sense.

In this example, the media server was previously invited to the conference indicated.

```
C->S: RECORD /meeting/audio.en RTSP/1.0 954
      Session: 1234
      Conference: 128.16.64.19/32492374
```

## 9.11 Embedded Binary Data

Binary packets such as RTP data are encapsulated by an ASCII dollar sign (24 decimal), followed by a one-byte session identifier, followed by the length of the encapsulated binary data as a binary, two-byte integer in network byte order. The binary data follows immediately afterwards, without a CRLF.

## 10 Status Code Definitions

Where applicable, HTTP status [H10] codes are re-used. Status codes that have the same meaning are not repeated here. See Table 1 for a listing of which status codes may be returned by which request.

### 10.1 Redirection 3xx

See [H10.3].

Within RTSP, redirection may be used for load balancing or redirecting stream requests to a server topologically closer to the client. Mechanisms to determine topological proximity are beyond the scope of this specification.

### 10.2 Client Error 4xx

#### 10.2.1 451 Parameter Not Understood

The recipient of the request does not support one or more parameters contained in the request.

#### 10.2.2 452 Conference Not Found

The conference indicated by a Conference header field is unknown to the media server.

#### 10.2.3 453 Not Enough Bandwidth

The request was refused since there was insufficient bandwidth. This may, for example, be the result of a resource reservation failure.

#### 10.2.4 45x Session Not Found

The RTSP session identifier is invalid or has timed out.

#### 10.2.5 45x Method Not Valid in This State

The client or server cannot process this request in its current state.

#### 10.2.6 45x Header Field Not Valid for Resource

The server could not act on a required request header. For example, if PLAY contains the Range header field, but the stream does not allow seeking.

#### 10.2.7 45x Invalid Range

The Range value given is out of bounds, e.g., beyond the end of the presentation.

#### 10.2.8 45x Parameter Is Read-Only

The parameter to be set by SET\_PARAMETER can only be read, but not modified.

### 11 Header Field Definitions

HTTP/1.1 or other, non-standard header fields not listed here currently have no well-defined meaning and SHOULD be ignored by the recipient.

Table 3 summarizes the header fields used by RTSP. Type "R" designates request headers, type "r" response headers. Fields marked with "req." in the column labeled "support" MUST be implemented by the recipient for a particular method, while fields marked "opt." are optional. Note that not all fields marked 'r' will be sent in every request of this type; merely, that client (for response headers) and server (for request headers) MUST implement them. The last column lists the method for which this header field is meaningful; the designation "entity" refers to all methods that return a message body. Within this specification, DESCRIBE and GET\_PARAMETER fall into this class.

If the field content does not apply to the particular resource, the server MUST return status 45x (Header Field Not Valid for Resource).

#### 11.1 Accept

The Accept request-header field can be used to specify certain presentation description content types which are acceptable for the response.

The "level" parameter for presentation descriptions is properly defined as part of the MIME type registration, not here.

See [H14.1] for syntax.

Example of use:

Accept: application/rtsp, application/sdp;level=2

Header	type	support	methods
Accept	R	opt.	entity
Accept-Encoding	R	opt.	entity
Accept-Language	R	opt.	all
Authorization	R	opt.	all
Bandwidth	R	opt.	SETUP
Blocksize	R	opt.	all but OPTIONS, TEARDOWN
Cache-Control	Rr	opt.	SETUP
Conference	R	opt.	SETUP
Connection	Rr	req.	all
Content-Encoding	R	req.	SET_PARAMETER
Content-Encoding	r	req.	DESCRIBE
Content-Length	R	req.	SET_PARAMETER
Content-Length	r	req.	entity
Content-Type	R	req.	SET_PARAMETER
Content-Type	r	req.	entity
Date	Rr	opt.	all
Expires	r	opt.	DESCRIBE
If-Modified-Since	R	opt.	DESCRIBE, SETUP
Last-Modified	r	opt.	entity
Public	r	opt.	all
Range	R	opt.	PLAY, PAUSE, RECORD
Range	r	opt.	PLAY, PAUSE, RECORD
Referer	R	opt.	all
Require	R	req.	all
Retry-After	r	opt.	all
Scale	Rr	opt.	PLAY, RECORD
Session	Rr	req.	all but SETUP, OPTIONS
Server	r	opt.	all
Speed	Rr	opt.	PLAY
Transport	Rr	req.	SETUP
Transport-Require	R	xeq.	all
User-Agent	R	opt.	all
Via	Rr	opt.	all
WWW-Authenticate	r	opt.	all

Table 3: Overview of RTSP header fields

## 11.2 Accept-Encoding

See [H14.3]

## 11.3 Accept-Language

See [H14.4]. Note that the language specified applies to the presentation description and any reason phrases, not the media



content.

#### 11.4 Allow

The Allow response header field lists the methods supported by the resource identified by the request-URI. The purpose of this field is to strictly inform the recipient of valid methods associated with the resource. An Allow header field must be present in a 405 (Method not allowed) response.

Example of use:

Allow: SETUP, PLAY, RECORD, SET\_PARAMETER

#### 11.5 Authorization

See [H14.8]

#### 11.6 Bandwidth

The Bandwidth request header field describes the estimated bandwidth available to the client, expressed as a positive integer and measured in bits per second.

Bandwidth = "Bandwidth" ":" 1\*DIGIT

Example:

Bandwidth: 4000

#### 11.7 Blocksize

This request header field is sent from the client to the media server asking the server for a particular media packet size. This packet size does not include lower-layer headers such as IP, UDP, or RTP. The server is free to use a blocksize which is lower than the one requested. The server MAY truncate this packet size to the closest multiple of the minimum media-specific block size or overrides it with the media specific size if necessary. The block size is a strictly positive decimal number and measured in octets. The server only returns an error (416) if the value is syntactically invalid.

## 11.8 Cache-Control

The Cache-Control general header field is used to specify directives that MUST be obeyed by all caching mechanisms along the request/response chain.

Cache directives must be passed through by a proxy or gateway application, regardless of their significance to that application, since the directives may be applicable to all recipients along the request/response chain. It is not possible to specify a cache-directive for a specific cache.

Cache-Control should only be specified in a SETUP request and its response. Note: Cache-Control does not govern the caching of responses as for HTTP, but rather of the stream identified by the SETUP request. Responses to RTSP requests are not cacheable.

[HS: Should there be an exception for DESCRIBE?]

```
Cache-Control    = "Cache-Control" ":" 1#cache-directive
```

```
cache-directive = cache-request-directive
                  | cache-response-directive
```

```
cache-request-directive =
    "no-cache"
    | "max-stale"
    | "min-fresh"
    | "only-if-cached"
    | cache-extension
```

```
cache-response-directive =
    "public"
    | "private"
    | "no-cache"
    | "no-transform"
    | "must-revalidate"
    | "proxy-revalidate"
    | "max-age" "=" delta-seconds
    | cache-extension
```

```
cache-extension = token [ "=" ( token | quoted-string ) ]
```

no-cache: Indicates that the media stream MUST NOT be cached anywhere. This allows an origin server to prevent caching even

by caches that have been configured to return stale responses to client requests.

**public:** Indicates that the media stream is cachable by any cache.

**private:** Indicates that the media stream is intended for a single user and **MUST NOT** be cached by a shared cache. A private (non-shared) cache may cache the media stream.

**no-transform:** An intermediate cache (proxy) may find it useful to convert the media type of certain stream. A proxy might, for example, convert between video formats to save cache space or to reduce the amount of traffic on a slow link. Serious operational problems may occur, however, when these transformations have been applied to streams intended for certain kinds of applications. For example, applications for medical imaging, scientific data analysis and those using end-to-end authentication, all depend on receiving a stream that is bit for bit identical to the original entity-body. Therefore, if a response includes the no-transform directive, an intermediate cache or proxy **MUST NOT** change the encoding of the stream. Unlike HTTP, RTSP does not provide for partial transformation at this point, e.g., allowing translation into a different language.

**only-if-cached:** In some cases, such as times of extremely poor network connectivity, a client may want a cache to return only those media streams that it currently has stored, and not to receive these from the origin server. To do this, the client may include the only-if-cached directive in a request. If it receives this directive, a cache **SHOULD** either respond using a cached media stream that is consistent with the other constraints of the request, or respond with a 504 (Gateway Timeout) status. However, if a group of caches is being operated as a unified system with good internal connectivity, such a request **MAY** be forwarded within that group of caches.

**max-stale:** Indicates that the client is willing to accept a media stream that has exceeded its expiration time. If max-stale is assigned a value, then the client is willing to accept a response that has exceeded its expiration time by no more than the specified number of seconds. If no value is assigned to max-stale, then the client is willing to accept a stale response of any age.

**min-fresh:** Indicates that the client is willing to accept a media stream whose freshness lifetime is no less than its current age plus the specified time in seconds. That is, the client wants a

response that will still be fresh for at least the specified number of seconds.

**must-revalidate:** When the must-revalidate directive is present in a SETUP response received by a cache, that cache MUST NOT use the entry after it becomes stale to respond to a subsequent request without first revalidating it with the origin server. (I.e., the cache must do an end-to-end revalidation every time, if, based solely on the origin server's Expires, the cached response is stale.)

### 11.9 Conference

This request header field establishes a logical connection between a conference, established using non-RTSP means, and an RTSP stream. The conference-id must not be changed for the same RTSP session.

Conference = "Conference" ":" conference-id

#### Example:

Conference: 199702170042.SAA08642@obiwan.arl.wustl.edu

### 11.10 Connection

See [H14.10].

### 11.11 Content-Encoding

See [H14.12]

### 11.12 Content-Length

This field contains the length of the content of the method (i.e. after the double CRLF following the last header). Unlike HTTP, it MUST be included in all messages that carry content beyond the header portion of the message. It is interpreted according to [H14.14].

### 11.13 Content-Type

See [H14.18]. Note that the content types suitable for RTSP are likely to be restricted in practice to presentation descriptions and parameter-value types.

## 11.14 Date

See [H14.19].

## 11.15 Expires

The Expires entity-header field gives the date/time after which the media-stream should be considered stale. A stale cache entry may not normally be returned by a cache (either a proxy cache or an user agent cache) unless it is first validated with the origin server (or with an intermediate cache that has a fresh copy of the entity). See section 13.2 for further discussion of the expiration model.

The presence of an Expires field does not imply that the original resource will change or cease to exist at, before, or after that time.

The format is an absolute date and time as defined by HTTP-date in [H3.3]; it MUST be in RFC1123-date format:

Expires = "Expires" ":" HTTP-date

An example of its use is

Expires: Thu, 01 Dec 1994 16:00:00 GMT

RTSP/1.0 clients and caches MUST treat other invalid date formats, especially including the value "0", as in the past (i.e., "already expired").

To mark a response as "already expired," an origin server should use an Expires date that is equal to the Date header value.

To mark a response as "never expires," an origin server should use an Expires date approximately one year from the time the response is sent. RTSP/1.0 servers should not send Expires dates more than one year in the future.

The presence of an Expires header field with a date value of some time in the future on a media stream that otherwise would by default be non-cacheable indicates that the media stream is cachable, unless indicated otherwise by a Cache-Control header field (Section 11.8).

#### 11.16 If-Modified-Since

The If-Modified-Since request-header field is used with the DESCRIBE and SETUP methods to make them conditional: if the requested variant has not been modified since the time specified in this field, a description will not be returned from the server ( DESCRIBE) or a stream will not be setup ( SETUP); instead, a 304 (not modified) response will be returned without any message-body.

If-Modified-Since = "If-Modified-Since" ":" HTTP-date

An example of the field is:

If-Modified-Since: Sat, 29 Oct 1994 19:43:31 GMT

#### 11.17 Last-modified

The Last-Modified entity-header field indicates the date and time at which the origin server believes the variant was last modified. See [H14.29]. If the request URI refers to an aggregate, the field indicates the last modification time across all leave nodes of that aggregate.

#### 11.18 Location

See [H14.30].

#### 11.19 Nack-Transport-Require

Negative acknowledgement of features not supported by the server. If there is a proxy on the path between the client and the server, the proxy MUST insert a message reply with an error message 506 (Feature not supported).

HS: Same caveat as for Require applies.

#### 11.20 Range

This request header field specifies a range of time. The range can be specified in a number of units. This specification defines the smpte (see Section 3.4) and clock (see Section 3.6) range units. Within

RTSP, byte ranges [H14.36.1] are not meaningful and MUST NOT be used. The header may also contain a time parameter in UTC, specifying the time at which the operation is to be made effective.

```
Range = "Range" ":" 1#ranges-specifier [ ";" "time" "=" utc-time ]  
ranges-specifier = npt-range | utc-range | smpte-range
```

Example:

```
Range: clock=19960213T143205Z-;Time=19970123T143720Z
```

The notation is similar to that used for the HTTP/1.1 header. It allows to select a clip from the media object, to play from a given point to the end and from the current location to a given point.

## 11.21 Require

The Require header is used by clients to query the server about features that it may or may not support. The server MUST respond to this header by negatively acknowledging those features which are NOT supported in the Unsupported header.

HS: Naming of features -- yet another name space. I believe this header field to be redundant. PEP should be used instead.

For example

```
C->S:  SETUP /foo/bar/baz.rm RTSP/1.0 302  
      Require: funky-feature  
      Funky-Parameter: funkystuff  
  
S->C:  RTSP/1.0 200 506 Option not supported  
      Unsupported: funky-feature  
  
C->S:  SETUP /foo/bar/baz.rm RTSP/1.0 303  
  
S->C:  RTSP/1.0 200 303 OK
```

This is to make sure that the client-server interaction will proceed optimally when all options are understood by both sides, and only slow down if options aren't understood (as in the case above). For a well-matched client-server pair, the interaction proceeds quickly, saving a round-trip often required by negotiation mechanisms. In addition, it also removes state ambiguity when the client requires features that the server doesn't understand.

#### 11.22 Retry-After

See [H14.38].

#### 11.23 Scale

A scale value of 1 indicates normal play or record at the normal forward viewing rate. If not 1, the value corresponds to the rate with respect to normal viewing rate. For example, a ratio of 2 indicates twice the normal viewing rate ("fast forward") and a ratio of 0.5 indicates half the normal viewing rate. In other words, a ratio of 2 has normal play time increase at twice the wallclock rate. For every second of elapsed (wallclock) time, 2 seconds of content will be delivered. A negative value indicates reverse direction.

Unless requested otherwise by the Speed parameter, the data rate SHOULD not be changed. Implementation of scale changes depends on the server and media type. For video, a server may, for example, deliver only key frames or selected key frames. For audio, it may time-scale the audio while preserving pitch or, less desirably, deliver fragments of audio.

The server should try to approximate the viewing rate, but may restrict the range of scale values that it supports. The response MUST contain the actual scale value chosen by the server.

If the request contains a Range parameter, the new scale value will take effect at that time.

Scale = "Scale" ":" [ "-" ] 1\*DIGIT [ "." \*DIGIT ]

Example of playing in reverse at 3.5 times normal rate:

Scale: -3.5



## 11.24 Speed

This request header fields parameter requests the server to deliver data to the client at a particular speed, contingent on the server's ability and desire to serve the media stream at the given speed. Implementation by the server is OPTIONAL. The default is the bit rate of the stream.

The parameter value is expressed as a decimal ratio, e.g., a value of 2.0 indicates that data is to be delivered twice as fast as normal. A speed of zero is invalid. A negative value indicates that the stream is to be played back in reverse direction.

HS: With 'Scale', the negative value is redundant and should probably be removed since it only leads to possible conflicts when Scale is positive and Speed negative.

If the request contains a Range parameter, the new speed value will take effect at that time.

Speed = "Speed" ":" [ "-" ] 1\*DIGIT [ "." \*DIGIT ]

Example:

Speed: 2.5

## 11.25 Server

See [H14.39]

## 11.26 Session

This request and response header field identifies an RTSP session, started by the media server in a SETUP response and concluded by TEARDOWN on the presentation URL. The session identifier is chosen by the media server and has the same syntax as a conference identifier. Once a client receives a Session identifier, it MUST return it for any request related to that session.

HS: This may be redundant with the standards-track HTTP state maintenance mechanism [2]. The equivalent way of

doing this would be for the server to send Set-Cookie: Session="123"; Version=1; Path = "/twister" and for the client to return later Cookie: Session = "123"; \$Version=1; \$Path = "/twister" response to the TEARDOWN message, the server would simply send Set-Cookie: Session="123"; Version=1; Max-Age=0 to get rid of the cookie on the client side. Cookies also have a time-out, so that a server may limit the lifetime of a session at will. Unlike a web browser, a client would not store these states on disk. To avoid privacy issues, we should prohibit the Host parameter.

## 11.27 Transport

This request header indicates which transport protocol is to be used and configures its parameters such as multicast, compression, multicast time-to-live and destination port for a single stream. It sets those values not already determined by a presentation description. In some cases, the presentation description contains all necessary information. In those cases, a Transport header field (and the SETUP request containing it) are not needed.

in whatever protocol is being used by the control stream. Currently, the next-layer protocols RTP is defined. Parameters may be added to each protocol, separated by a semicolon. For RTP, the boolean parameter compressed is defined, indicating compressed RTP according to RFC XXXX. For multicast UDP, the integer parameter ttl defines the time-to-live value to be used. The client may specify the multicast address with the multicast parameter. A server SHOULD authenticate the client before allowing the client to direct a media stream to a multicast address not chosen by the server to avoid becoming the unwitting perpetrator of a denial-of-service attack. For UDP and TCP, the parameter port defines the port data is to be sent to.

The SSRC parameter indicates the RTP SSRC value that should be (request) or will be (response) used by the media server. This parameter is only valid for unicast transmission. It identifies the synchronization source to be associated with the media stream.

The Transport header MAY also be used to change certain transport parameters. A server MAY refuse to change parameters of an existing stream.

The server MAY return a Transport response header in the response to indicate the values actually chosen.

A Transport request header field may contain a list of transport

options acceptable to the client. In that case, the server MUST return a single option which was actually chosen. The Transport header field makes sense only for an individual media stream, not a presentation.

```

Transport = "Transport" ":"
           1#transport-protocol/upper-layer *parameter
transport-protocol = "UDP" | "TCP"
upper-layer       = "RTP"
parameters = ";" "multicast" [ "=" mca ]
               | ";" "compressed"
               | ";" "interleaved"
               | ";" "ttl" "=" ttl
               | ";" "port" "=" port
               | ";" "ssrc" "=" ssrc
ttl             = 1*3(DIGIT)
port            = 1*5(DIGIT)
ssrc            = 8*8(HEX)
mca             = host

```

Example:

```
Transport: udp/rtp;compressed;ttl=127;port=3456
```

## 11.28 Transport-Require

The Transport-Require header is used to indicate proxy-sensitive features that MUST be stripped by the proxy to the server if not supported. Furthermore, any Transport-Require header features that are not supported by the proxy MUST be negatively acknowledged by the proxy to the client if not supported.

See Section 11.21 for more details on the mechanics of this message and a usage example.

HS: Same caveat as for Require applies.

## 11.29 Unsupported

See Section 11.21 for a usage example.

HS: same caveat as for Require applies.

## 11.30 User-Agent

See [H14.42]

## 11.31 Via

See [H14.44].

## 11.32 WWW-Authenticate

See [H14.46].

## 12 Caching

In HTTP, response-request pairs are cached. RTSP differs significantly in that respect. Responses are not cachable, with the exception of the stream description returned by DESCRIBE. (Since the responses for anything but DESCRIBE and GET\_PARAMETER do not return any data, caching is not really an issue for these requests.) However, it is desirable for the continuous media data, typically delivered out-of-band with respect to RTSP, to be cached.

On receiving a SETUP or PLAY request, the proxy would ascertain as to whether it has an up-to-date copy of the continuous media content. If not, it would modify the SETUP transport parameters as appropriate and forward the request to the origin server. Subsequent control commands such as PLAY or PAUSE would pass the proxy unmodified. The proxy would then pass the continuous media data to the client, while possibly making a local copy for later re-use. The exact behavior allowed to the cache is given by the cache-response directives described in Section 11.8. A cache MUST answer any DESCRIBE requests if it is currently serving the stream to the requestor, as it is possible that low-level details of the stream description may have changed on the origin-server.

Note that an RTSP cache, unlike the HTTP cache, is of the "cut-through" variety. Rather than retrieving the whole resource from the origin server, the cache simply copies the streaming data as it passes by on its way to the client, thus, it does not introduce additional latency.

To the client, an RTSP proxy cache would appear like a regular media server, to the media origin server like a client. Just like an HTTP cache has to store the content type, content language, etc. for the objects it caches, a media cache has to store the presentation description. Typically, a cache would eliminate all transport-references (that is, multicast information) from the presentation description, since these are independent of the data delivery from

the cache to the client. Information on the encodings remains the same. If the cache is able to translate the cached media data, it would create a new presentation description with all the encoding possibilities it can offer.

### 13 Examples

The following examples reference stream description formats that are not finalized, such as RTSL and SDP. Please do not use these examples as a reference for those formats.

#### 13.1 Media on Demand (Unicast)

Client C requests a movie from media servers A ( audio.example.com ) and V ( video.example.com ). The media description is stored on a web server W. The media description contains descriptions of the presentation and all its streams, including the codecs that are available, dynamic RTP payload types, the protocol stack and content information such as language or copyright restrictions. It may also give an indication about the time line of the movie.

In our example, the client is only interested in the last part of the movie. The server requires authentication for this movie. The audio track can be dynamically switched between between two sets of encodings. The URL with scheme rtspu indicates the media servers want to use UDP for exchanging RTSP messages.

```
C->W: DESCRIBE /twister HTTP/1.1
Host: www.example.com
Accept: application/rtsl; application/sdp
```

```
W->C: 200 OK
Content-Type: application/rtsl
```

```
<session>
  <group language=en lipsync>
    <switch>
      <track type=audio
        e="PCMU/8000/1"
        src="rtsp://audio.example.com/twister/audio.en/lofi">
      <track type=audio
        e="DVI4/16000/2" pt="90 DVI4/8000/1"
        src="rtsp://audio.example.com/twister/audio.en/hifi">
    </switch>
    <track type="video/jpeg"
      src="rtspu://video.example.com/twister/video">
  </group>
```

</session>

C->A: SETUP rtsp://audio.example.com/twister/audio.en/lofi RTSP/1.0 1  
Transport: rtp/udp;compression;port=3056

A->C: RTSP/1.0 200 1 OK  
Session: 1234

C->V: SETUP rtsp://video.example.com/twister/video RTSP/1.0 1  
Transport: rtp/udp;compression;port=3058

V->C: RTSP/1.0 200 1 OK  
Session: 1235

C->V: PLAY rtsp://video.example.com/twister/video RTSP/1.0 2  
Session: 1235  
Range: smpte=0:10:00-

V->C: RTSP/1.0 200 2 OK

C->A: PLAY rtsp://audio.example.com/twister/audio.en/lofi RTSP/1.0 2  
Session: 1234  
Range: smpte=0:10:00-

A->C: 200 2 OK

C->A: TEARDOWN rtsp://audio.example.com/twister/audio.en/lofi RTSP/1.0 3  
Session: 1234

A->C: 200 3 OK

C->V: TEARDOWN rtsp://video.example.com/twister/video RTSP/1.0 3  
Session: 1235

V->C: 200 3 OK

Even though the audio and video track are on two different servers, may start at slightly different times and may drift with respect to each other, the client can synchronize the two using standard RTP methods, in particular the time scale contained in the RTCP sender reports.

### 13.2 Live Media Presentation Using Multicast

The media server M chooses the multicast address and port. Here, we assume that the web server only contains a pointer to the full

description, while the media server M maintains the full description. During the RTSP session, a new subtitling stream is added.

```
C->W: GET /concert HTTP/1.1
      Host: www.example.com

W->C: HTTP/1.1 200 OK
      Content-Type: application/rtsl

      <session>
        <track id=17 src="rtsp://live.example.com/concert/audio">
      </session>

C->M: DESCRIBE rtsp://live.example.com/concert/audio RTSP/1.0 1

M->C: RTSP/1.0 200 1 OK
      Content-Type: application/rtsl

      <track id=17 type=audio address=224.2.0.1 port=3456 ttl=16>

C->M: SETUP rtsp://live.example.com/concert/audio RTSP/1.0 2
      Transport: multicast=224.2.0.1; port=3456; ttl=16

C->M: PLAY rtsp://live.example.com/concert/audio RTSP/1.0 3
      Range: smpte 1:12:0

M->C: RTSP/1.0 405 3 No positioning possible

M->C: DESCRIBE concert RTSP/1.0
      Content-Type: application/rtsl

      <session>
        <track id=17
          media=audio/g.728 src="rtsp://live.example.com/concert/audio">
        <track id=18
          media=text/html src="rtsp://live.example.com/concert/lyrics">
      </session>

C->M: PLAY rtsp://live.example.com/concert/lyrics RTSP/1.0
```

The attempt to position the stream fails since this is a live presentation.

### 13.3 Playing media into an existing session

A conference participant C wants to have the media server M play back a demo tape into an existing conference. When retrieving the presentation description, C indicates to the media server that the network addresses and encryption keys are already given by the conference, so they should not be chosen by the server. The example omits the simple ACK responses.

```
C->M: GET /demo HTTP/1.1
      Host: www.example.com
      Accept: application/rtsp, application/sdp

M->C: HTTP/1.1 200 1 OK
      Content-type: application/rtsp

      <session>
        <track type=audio/g.723.1
          src="rtsp://server.example.com/demo/548/sound">
        </session>

C->M: SETUP rtsp://server.example.com/demo/548/sound RTSP/1.0 2
      Conference: 218kadjk
```

#### 13.4 Recording

The conference participant C asks the media server M to record a meeting. If the presentation description contains any alternatives, the server records them all.

```
C->M: DESCRIBE rtsp://server.example.com/meeting RTSP/1.0 89
      Content-Type: application/sdp

      v=0
      s=Mbone Audio
      i=Discussion of Mbone Engineering Issues

M->C: 415 89 Unsupported Media Type
      Accept: application/rtsp

C->M: DESCRIBE rtsp://server.example.com/meeting RTSP/1.0 90
      Content-Type: application/rtsp

M->C: 200 90 OK

C->M: RECORD rtsp://server.example.com/meeting RTSP/1.0 91
```



Range: clock 19961110T1925-19961110T2015

## 14 Syntax

The RTSP syntax is described in an augmented Backus-Naur form (BNF) as used in RFC 2068 (HTTP/1.1).

### 14.1 Base Syntax

```

OCTET      = <any 8-bit sequence of data>
CHAR       = <any US-ASCII character (octets 0 - 127)>
UPALPHA    = <any US-ASCII uppercase letter "A".."Z">
LOALPHA    = <any US-ASCII lowercase letter "a".."z">
ALPHA      = UPALPHA | LOALPHA
DIGIT      = <any US-ASCII digit "0".."9">
CTL        = <any US-ASCII control character
              (octets 0 - 31) and DEL (127)>
CR         = <US-ASCII CR, carriage return (13)>
LF         = <US-ASCII LF, linefeed (10)>
SP         = <US-ASCII SP, space (32)>
HT         = <US-ASCII HT, horizontal-tab (9)>
<">       = <US-ASCII double-quote mark (34)>
CRLF      = CR LF
LWS        = [CRLF] 1*( SP | HT )
TEXT       = <any OCTET except CTLs>
tspecials  = "(" | ")" | "<" | ">" | "@"
              | "," | ";" | ":" | "
              | "/" | "[" | "]" | "?" | "="
              | "{" | "}" | SP | HT
token      = 1*<any CHAR except CTLs or tspecials>
quoted-string = ( <"> *(qdtext) <"> )
qdtext     = <any TEXT except <">>
quoted-pair = "

message-header = field-name ":" [ field-value ] CRLF
field-name    = token
field-value   = *( field-content | LWS )
field-content = <the OCTETs making up the field-value and consisting
                  of either *TEXT or combinations of token, tspecials,
                  and quoted-string>

```

## 15 Security Considerations

The protocol offers the opportunity for a remote-control denial-of-service attack. The attacker, using a forged source IP address, can ask for a stream to be played back to that forged IP address.

Since there is no relation between a transport layer connection and an RTSP session, it is possible for a malicious client to issue requests with random session identifiers which would affect unsuspecting clients. This does not require spoofing of network packet addresses. The server SHOULD use a large random session identifier to make this attack more difficult.

Both problems can be prevented by appropriate authentication.

In addition, the security considerations outlined in [H15] apply.

#### A RTSP Protocol State Machines

The RTSP client and server state machines describe the behavior of the protocol from RTSP session initialization through RTSP session termination.

[TBD: should we allow for the trivial case of a server that only implements the PLAY message, with no control.]

State is defined on a per object basis. An object is uniquely identified by the stream URL and the RTSP session identifier. (A server may choose to generate dynamic presentation descriptions where the URL is unique for a particular RTSP session and thus may not need an explicit RTSP session identifier in the request header.) Any request/reply using URLs denoting an RTSP session comprised of multiple streams will have an effect on the individual states of all the substreams. For example, if the stream /movie contains two substreams /movie/audio and /movie/video, then the following command:

```
PLAY /movie RTSP/1.0 559
Session: 12345
```

will have an effect on the states of movie/audio and movie/video.

This example does not imply a standard way to represent substreams in URLs or a relation to the filesystem. See Section 3.2.

The requests OPTIONS, DESCRIBE, GET\_PARAMETER, SET\_PARAMETER do

not have any effect on client or server state and are therefore not listed in the state tables.

Client and servers MUST disregard messages with a sequence number less than the last one. If no message has been received, the first received message's sequence number will be the starting point.

A.1 Client State Machine

The client can assume the following states:

Init: SETUP has been sent, waiting for reply.

Ready: SETUP reply received OR after playing, PAUSE reply received.

Playing: PLAY reply received

Recording: RECORD reply received

The client changes state on receipt of replies to requests. If no explicit SETUP is required for the object (for example, it is available via a multicast group), state begins at READY. In this case, there are only two states, READY and PLAYING.

The "next state" column indicates the state assumed after receiving a success response (2xx). If a request yields a status code greater or equal to 300, the client state becomes Init, with the exception of status codes 401 (Unauthorized) and 402 (Payment Required), where the state remains unchanged and the request should be re-issued with the appropriate authentication or payment information. Messages not listed for each state MUST NOT be issued by the client in that state, with the exception of messages not affecting state, as listed above. Receiving a REDIRECT from the server is equivalent to receiving a 3xx redirect status from the server.

HS: Depends on allowing PLAY without SETUP. After 4xx or 5xx error, do we go back to Init?

state	message	next state
Init	SETUP	Ready
	TEARDOWN	Init
Ready	PLAY	Playing
	RECORD	Recording
Playing	TEARDOWN	Init
	PAUSE	Ready
	TEARDOWN	Init

Recording	PLAY	Playing
	RECORD	Recording
	SETUP	Playing (changed transport)
	PAUSE	Init
	TEARDOWN	Init
	PLAY	Playing
	RECORD	Recording
	SETUP	Recording (changed transport)

## A.2 Server State Machine

The server can assume the following states:

Init: The initial state, no valid SETUP received.

Ready: Last SETUP received was successful, reply sent or after playing, last PAUSE received was successful, reply sent.

Playing: Last PLAY received was successful, reply sent. Data is being sent.

Recording: The server is recording media data.

The server changes state on receiving requests. If the server is in state Playing or Recording and in unicast mode, it MAY revert to Init and tear down the RTSP session if it has not received "wellness" information, such as RTCP reports, from the client for a defined interval, with a default of one minute. If the server is in state Ready, it MAY revert to Init if it does not receive an RTSP request for an interval of more than one minute.

The REDIRECT message, when sent, is effective immediately. If a similar change of location occurs at a certain time in the future, this is assumed to be indicated by the presentation description.

SETUP is valid in states Init and Ready only. An error message should be returned in other cases. If no explicit SETUP is required for the object, state starts at READY, there are only two states READY and PLAYING.

state	message	next state
Init	SETUP	Ready
	TEARDOWN	Init
Ready	PLAY	Playing
	SETUP	Ready
	TEARDOWN	Ready

Playing	PLAY	Playing
	PAUSE	Ready
	TEARDOWN	Ready
	RECORD	Recording
	SETUP	Playing
Recording	RECORD	Recording
	PAUSE	Ready
	TEARDOWN	Ready
	PLAY	Playing
	SETUP	Recording

## B Open Issues

- o Define text/rtsp-parameter MIME type.
- o HS believes that RTSP should only control individual media objects rather than aggregates. This avoids disconnects between presentation descriptions and streams and avoids having to deal separately with single-host and multi-host case. Cost: several PLAY/PAUSE/RECORD in one packet, one for each stream.
- o Allow changing of transport for a stream that's playing? May not be a great idea since the same can be accomplished by tear down and re-setup.
- o Allow fragment (#) identifiers for controlling substreams in Quicktime, AVI and ASF files?
- o How does the server get back to the client unless a persistent connection is used? Probably cannot, in general.
- o Cache and proxy behavior?
- o Session: or Set-Cookie: ?
- o When do relative RTSP URLs make sense?
- o Nack-require, etc. are dubious. This is getting awfully close to the HTTP extension mechanisms [19] in complexity, but is different.
- o Use HTTP absolute path + Host field or do the right thing and carry full URL, including host in request?

## C Changes

Since the February 1997 version, the following changes were made:

- o Various editorial changes and clarifications.
- o Removed references to SDF and replaced by RTSL.
- o Added Scale general header.
- o Clarify behavior of PLAY.
- o Rename GET to DESCRIBE.
- o Removed SESSION since it is just DESCRIBE in the other direction.
- o Rename CLOSE to TEARDOWN, in symmetry with SETUP.
- o Terminology adjusted to "presentation" and "stream".
- o Redundant syntax BNF in appendix removed since it just duplicates HTTP spec.
- o Beginnings of cache control.

Changes are marked by changebars in the margins of the PostScript version.

#### D Author Addresses

Henning Schulzrinne  
Dept. of Computer Science  
Columbia University  
1214 Amsterdam Avenue  
New York, NY 10027  
USA  
electronic mail: schulzrinne@cs.columbia.edu

Anup Rao  
Netscape Communications Corp.  
USA  
electronic mail: anup@netscape.com

Robert Lanphier  
Progressive Networks  
1111 Third Avenue Suite 2900  
Seattle, WA 98101  
USA  
electronic mail: robila@prognnet.com

#### E Acknowledgements

H. Schulzrinne, A. Rao, R. Lanphier

[Page 57]

This draft is based on the functionality of the RTSP draft. It also borrows format and descriptions from HTTP/1.1.

This document has benefited greatly from the comments of all those participating in the MMUSIC-WG. In addition to those already mentioned, the following individuals have contributed to this specification:

Rahul Agarwal	Eduardo F. Llach
Bruce Butterfield	Rob McCool
Martin Dunsmuir	Sujal Patel
Eric Fleischman	
Mark Handley	Igor Plotnikov
Peter Haight	Pinaki Shah
Brad Hefta-Gaub	Jeff Smith
John K. Ho	Alexander Sokolsky
Ruth Lang	Dale Stammen
Stephanie Leif	John Francis Stracke

## F Bibliography

- [1] H. Schulzrinne, "RTP profile for audio and video conferences with minimal control," RFC 1890, Internet Engineering Task Force, Jan. 1996.
- [2] D. Kristol and L. Montulli, "HTTP state management mechanism," RFC 2109, Internet Engineering Task Force, Feb. 1997.
- [3] F. Yergeau, G. Nicol, G. Adams, and M. Duerst, "Internationalization of the hypertext markup language," RFC 2070, Internet Engineering Task Force, Jan. 1997.
- [4] S. Bradner, "Key words for use in RFCs to indicate requirement levels," Internet Draft, Internet Engineering Task Force, Jan. 1997. Work in progress.
- [5] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, and T. Berners-Lee, "Hypertext transfer protocol -- HTTP/1.1," RFC 2068, Internet Engineering Task Force, Jan. 1997.
- [6] M. Handley, "SDP: Session description protocol," Internet Draft, Internet Engineering Task Force, Nov. 1996. Work in progress.
- [7] A. Freier, P. Karlton, and P. Kocher, "The TLS protocol," Internet Draft, Internet Engineering Task Force, Dec. 1996. Work in progress.

[8] J. Franks, P. Hallam-Baker, J. Hostetler, P. A. Luotonen, and E. L. Stewart, "An extension to HTTP: digest access authentication," RFC 2069, Internet Engineering Task Force, Jan. 1997.

[9] J. Postel, "User datagram protocol," STD 6, RFC 768, Internet Engineering Task Force, Aug. 1980.

[10] R. Hinden and C. Partridge, "Version 2 of the reliable data protocol (RDP)," RFC 1151, Internet Engineering Task Force, Apr. 1990.

[11] J. Postel, "Transmission control protocol," STD 7, RFC 793, Internet Engineering Task Force, Sept. 1981.

[12] M. Handley, H. Schulzrinne, and E. Schooler, "SIP: Session initiation protocol," Internet Draft, Internet Engineering Task Force, Dec. 1996. Work in progress.

[13] P. McMahon, "GSS-API authentication method for SOCKS version 5," RFC 1961, Internet Engineering Task Force, June 1996.

[14] D. Crocker, "Augmented BNF for syntax specifications: ABNF," Internet Draft, Internet Engineering Task Force, Oct. 1996. Work in progress.

[15] R. Elz, "A compact representation of IPv6 addresses," RFC 1924, Internet Engineering Task Force, Apr. 1996.

[16] T. Berners-Lee, L. Masinter, and M. McCahill, "Uniform resource locators (URL)," RFC 1738, Internet Engineering Task Force, Dec. 1994.

[17] International Telecommunication Union, "Visual telephone systems and equipment for local area networks which provide a non-guaranteed quality of service," Recommendation H.323, Telecommunication Standardization Sector of ITU, Geneva, Switzerland, May 1996.

[18] ISO/IEC, "Information technology -- generic coding of moving pictures and associated audio informaiton -- part 6: extension for digital storage media and control," Draft International Standard ISO 13818-6, International Organization for Standardization ISO/IEC JTC1/SC29/WG11, Geneva, Switzerland, Nov. 1995.

[19] D. Connolly, "PEP: an extension mechanism for http," Internet Draft, Internet Engineering Task Force, Jan. 1997. Work in progress.



Table of Contents

1 Introduction ..... 1

1.1 Purpose ..... 1

1.2 Requirements ..... 3

1.3 Terminology ..... 3

1.4 Protocol Properties ..... 5

1.5 Extending RTSP ..... 6

1.6 Overall Operation ..... 7

1.7 RTSP States ..... 8

1.8 Relationship with Other Protocols ..... 9

2 Notational Conventions ..... 10

3 Protocol Parameters ..... 10

3.1 H3.1 ..... 10

3.2 RTSP URL ..... 10

3.3 Conference Identifiers ..... 11

3.4 SMPTE Relative Timestamps ..... 12

3.5 Normal Play Time ..... 13

3.6 Absolute Time ..... 13

4 RTSP Message ..... 13

4.1 Message Types ..... 14

4.2 Message Headers ..... 14

4.3 Message Body ..... 14

4.4 Message Length ..... 14

5 Request ..... 15

6 Response ..... 16

6.1 Status-Line ..... 17

6.1.1 Status Code and Reason Phrase ..... 17

6.1.2 Response Header Fields ..... 19

7 Entity ..... 19

7.1 Entity Header Fields ..... 21

7.2 Entity Body ..... 21

8 Connections ..... 21

8.1 Pipelining ..... 22

8.2 Reliability and Acknowledgements ..... 22

9 Method Definitions ..... 23

9.1 OPTIONS ..... 24

9.2 DESCRIBE ..... 25

9.3 SETUP ..... 26

9.4 PLAY ..... 27

9.5 PAUSE ..... 28

9.6 TEARDOWN ..... 30

9.7 GET\_PARAMETER ..... 30

9.8 SET\_PARAMETER ..... 31

9.9 REDIRECT ..... 31

9.10 RECORD ..... 32

9.11	Embedded Binary Data .....	32
10	Status Code Definitions .....	33
10.1	Redirection 3xx .....	33
10.2	Client Error 4xx .....	33
10.2.1	451 Parameter Not Understood .....	33
10.2.2	452 Conference Not Found .....	33
10.2.3	453 Not Enough Bandwidth .....	33
10.2.4	45x Session Not Found .....	33
10.2.5	45x Method Not Valid in This State .....	33
10.2.6	45x Header Field Not Valid for Resource .....	33
10.2.7	45x Invalid Range .....	33
10.2.8	45x Parameter Is Read-Only .....	34
11	Header Field Definitions .....	34
11.1	Accept .....	34
11.2	Accept-Encoding .....	35
11.3	Accept-Language .....	35
11.4	Allow .....	36
11.5	Authorization .....	36
11.6	Bandwidth .....	36
11.7	Blocksize .....	36
11.8	Cache-Control .....	37
11.9	Conference .....	39
11.10	Connection .....	39
11.11	Content-Encoding .....	39
11.12	Content-Length .....	39
11.13	Content-Type .....	39
11.14	Date .....	40
11.15	Expires .....	40
11.16	If-Modified-Since .....	41
11.17	Last-modified .....	41
11.18	Location .....	41
11.19	Nack-Transport-Require .....	41
11.20	Range .....	41
11.21	Require .....	42
11.22	Retry-After .....	43
11.23	Scale .....	43
11.24	Speed .....	44
11.25	Server .....	44
11.26	Session .....	44
11.27	Transport .....	45
11.28	Transport-Require .....	46
11.29	Unsupported .....	46
11.30	User-Agent .....	47
11.31	Via .....	47
11.32	WWW-Authenticate .....	47
12	Caching .....	47
13	Examples .....	48
13.1	Media on Demand (Unicast) .....	48

13.2 Live Media Presentation Using Multicast ..... 49

13.3 Playing media into an existing session ..... 50

13.4 Recording ..... 51

14 Syntax ..... 52

14.1 Base Syntax ..... 52

15 Security Considerations ..... 52

A RTSP Protocol State Machines ..... 53

A.1 Client State Machine ..... 54

A.2 Server State Machine ..... 55

B Open Issues ..... 56

C Changes ..... 56

D Author Addresses ..... 57

E Acknowledgements ..... 57

F Bibliography ..... 58